CS 6840 Algorithmic Game Theory November 18, 2024

Lecture 34: Games w/ collaboration and strong PoA

Instructor: Eva Tardos *Scribe: Eric Yachbes*

1 Introduction to Strong Nash and Collaboration

Last time, we covered equilibrium analysis. Today we will cover strong nash / collaboration. Collaborative game theory is its own large field, but strong nash is similar enough to our regular nash to cover.

Nash has an issue where everyone selfishly optimizes without talking to other players. If allowed to talk to players, you can "collaborate" for better results.

Game: Consider we have a game with n players, each have strategies s_i.

Recall that in regular nash, we want no player to be able to individually switch to a better strategy. Similarly, in strong nash, we want no subset of people to be able to all cooperate by all switching to different strategies to achieve a better result. Specifically, no player in this subset will be worse after all switch to this cooperative strategy. So, the difference between regular nash is now that a subset of people can switch strategies together and all of them achieve no worse of a result.

1.1 Definition of Strong Nash

Definition Formally, s is **not** a strong nash equilibrium iff there exists a subset A of players and a "collaborative strategy" s' s.t. no player switching to this collaborative strategy is worse:

$$c_i(s'_A, s_{-A}) \leq c_i(s)$$

for all $i \in A$ (the notation $c_i(s'_A, s_{-A})$ means we are want the cost for player i when all players $j \in A$ play strategy s'_j and all other players $j \notin A$ play strategy s_j) **and** there exists a player $j \in A$ that is strictly better off by switching to s':

$$c_j(s'_A, s_{-A}) < c_j(s)$$

Otherwise, s **is** a strong nash equilibrium. If we restrict $|A| = 1$, this is the same definition as regular nash. Thus, a strong nash is a regular nash. So the set of strong nash equilibrium is a subset of the set of regular nash equilbrium.

1.2 Strong Nash not guaranteed to exist

The issue is that a strong nash (even a mixed strong nash) is not guaranteed to exist.

Example: Prisoner's dilema does not have a strong nash. We have the following utilities:

	Defect	Cooperate
Defect	1,1	10,0
Cooperate	0,10	9,9

Any person would want to switch from cooperating, since they would be better off defecting. If both players are defecting, then both can switch to cooperating (this is not possible in regular nash, but it is possible here, since we allow collaboration). Thus, no strategy is a strong nash, since we can either greedily switch to defecting or cooperate together for a better result than both defecting.

1.3 Relation to Coalition Proof

Econ Concept: Strong Nash is stronger, but similar, to coalition proof[1]. A strategy is called coalition proof iff no subset of players can deviate to some other strategy that will be no worse for each player AND no subset of this deviating subset will want to switch strategies again. Essentially, we are allowed to collaborate once and then collaborate again. Thus, coalition proof only guards against deviation where no subset of the deviating subset has the incentive to deviate again. Thus it guards against fewer forms of deviation, making it weaker than strong nash.

For example, in prisoners dilema, there is no strong nash because we can deviate to both players collaborating. But, Prisoner's dillema is coalition proof because after deviating to both players collaborating, the players can deviate again to defecting as in regular nash.

2 Strong PoA in cost-sharing games

2.1 Definition of cost-sharing game

Recall: Players choose paths from source to sink on some graph. A set of playale strategy is $S_i \subseteq 2^E$. Each edge has some non-negative cost $c_e \geq 0$. x_e is the number of players that choose edge e in their path. The cost for player i is $c_i(s) = \sum_{e \in s_i} \frac{c_e}{x_e}$.

2.2 Existence of a regular nash with high PoA

We have a "bad" example in the following graph:

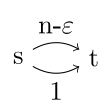

If we have everyone using the top edge, they will all have cost $\frac{n-\varepsilon}{n} = 1 - \frac{\varepsilon}{n}$. No individual will want to switch, since they will incur cost $1 - (1 - \frac{\varepsilon}{n}) = \frac{\varepsilon}{n}$. So, this is a nash.

But this is **not** a strong nash, since a group can collaborate to switch. For example, if two people agree to both switch to the bottom edge, then we have that they will get a new cost of 1/2, which can be less than their initial cost of $1 - \frac{\varepsilon}{n}$.

2.3 Proof of Strong PoA for general cost-sharing games

We can achieve a nicer PoA for strong nash equilibrium than for our regular nash PoA.

[1]Thanks to Ruqing Xu for help with writing the coalition proof section.

Theorem: PoA for strong nash $\leq H_n$ (harmonic number)

$$H_n = \sum_{i=1}^{n} 1/i$$

Proof: Say we have a strong nash s which is a vector $s_1, s_2, \ldots, s_n$ of strategies for each player. There is also an optimum strategy vector $s^* = s_1^*, s_2^*, \ldots, s_n^*$. What if everyone collaborates by deciding to all switch to s^*? The whole set would not like to switch (some person would be no better after switching to optimum), since s is a strong nash.

That is, there exists a player n (we let ourselves define the ordering arbitrarily in this way) s.t.

$$c_n(s^*) \geq c_n(s)$$

Now, define $A_{n-1} = \{1, 2, \ldots, n-1\}$ and say that this subset of players is collaboratively deciding to all switch to the optimum strategy s^*. Someone, say player $n-1$, out of this group would not be better off switching to the optimum, since s is a strong nash. That is, we have

$$c_{n-1}(s^*_{A_{n-1}}, s_{-A_{n-1}}) \geq c_{n-1}(s)$$

We continue this argument to get that for any i, with our chosen ordering of players, we have

$$c_i(s^*_{A_i}, s_{-A_i}) \geq c_i(s)$$

where $A_i = \{1, 2, \ldots, i\}$.

We will add up these inequalities. Notice that, since this is a cost-sharing game, the effect of players $-A_i$ can only decrease the cost c_i for player i. If we ignore this benefit, we get a looser bound

$$c_i(s^*_{A_i}) \geq c_i(s^*_{A_i}, s_{-A_i}) \geq c_i(s)$$

where $c_i(s^*_{A_i})$ is defined s.t. the players $-A_i$ do nothing. That is, the strategies are $s_1^*, \ldots, s_i^*, \emptyset, \ldots, \emptyset$, where $\emptyset$ represents not choosing any edges.

Summing up these inequalities, we get

$$cost(s) = \sum_i c_i(s) \leq \sum_i c_i(s^*_{A_i})$$

What can we say about $c_i(s^*_{A_i})$? Recall cost-sharing is a potential game, so there exists a function ϕ s.t.

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

We can write $c_i(s^*_{A_i}) = c_i(s^*_{A_i}) - c_i(s^*_{A_{i-1}}) = c_i(s^*_{A_i}) - 0$. This is because $c_i(s^*_{A_{i-1}})$ is an action where player i doesn't choose an edge, and so they incur no cost.

We now have

$$c_i(s^*_{A_i}) = \phi(s^*_{A_i}) - \phi(s^*_{A_{i-1}})$$

Now, when we sum up the inequalities, we get a telescoping sum:

$$\sum_i c_i(s^*_{A_i}) = \sum_i (\phi(s^*_{A_i}) - \phi(s^*_{A_{i-1}})) = \phi(s^*_{A_n}) - \phi(^*_{A_0}) = \phi(s^*) - \phi(\emptyset) = \phi(s^*)$$

Recall that we defined

$$\phi(s) = \sum_{\substack{e \in E, \\ x_e > 0}} \sum_{k=1}^{x_e} \frac{c_e}{k} = \sum_{\substack{e \in E, \\ x_e > 0}} H_{x_e} c_e$$

This implies $\phi(\emptyset) = 0$, since we sum over 0 edges.

Now,

$$cost(s^*) = \sum_i \sum_{e \in s_i^*} \frac{c_e}{x_e} = \sum_{\substack{e \in E, \\ x_e^* > 0}} \sum_{i=1}^{x_e} \frac{c_e}{x_e} = \sum_{\substack{e \in E, \\ x_e^* > 0}} c_e \le \phi(s^*) \le H_n cost(s^*)$$

(since $1 \le H_i \le H_n$ for all $1 \le i \le n$).

So, we can plug in the results from before to get

$$cost(s) = \sum_i c_i(s) \le \sum_i c_i(s^*_{A_i}) = \phi(s^*) \le H_n cost(s^*)$$

for strong nash s, proving that our strong PoA is $\le H_n$.

2.4 Bound on strong PoA in cost-sharing games is tight

Recall the example[2] from Lecture 20.

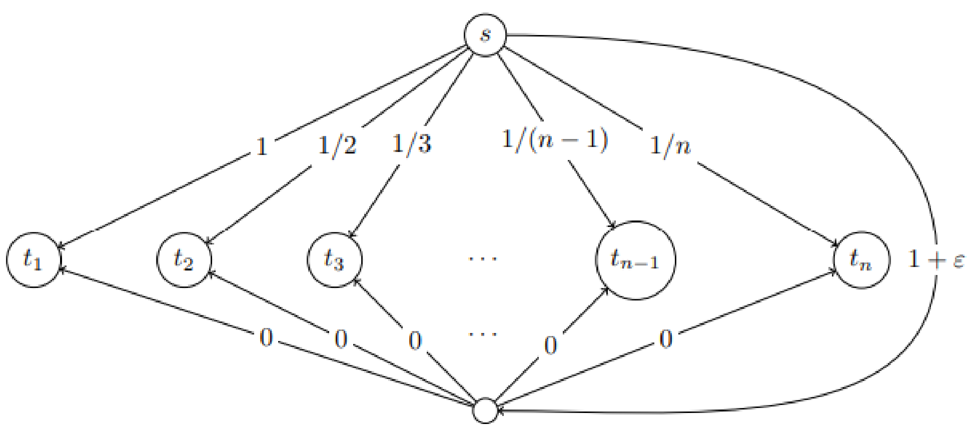

The optimal solution is for all players to take the $1 + \varepsilon$ edge. They will all incur cost $\frac{1+\varepsilon}{n}$, for a total cost of $1 + \varepsilon$.

However, player n would rather take the $1/n$ edge, since $1/n < \frac{1+\varepsilon}{n} = 1/n + \frac{\varepsilon}{n}$. Then player $n-1$ would rather take the $1/(n-1)$ edge, since $1/(n-1) < \frac{1+\varepsilon}{n-1} = 1/(n-1) + \frac{\varepsilon}{n-1}$, and so on. So everyone subsequently objects from taking the $1+\varepsilon$ edge in any nash. This means the only nash is when all players choose their own individual edge, giving a cost of $1 + 1/2 + \cdots + 1/n = H_n$.

This nash is also a strong nash, since if k people divert to the $1 + \varepsilon$ edge, then one of those people will have originally had cost $\le 1/k$ and their new cost will be $\frac{1+\varepsilon}{k} = 1/k + \varepsilon/k$, so they will not switch.

[2]Credit to Feras Al Taha for the picture

Since this is the only strong nash, and its cost is H_n, the PoA is $\frac{H_n}{1+\varepsilon}$, which can be arbitrarily close to H_n, showing our bound is tight. This example has Nash = Strong Nash (black solutions).

3 Concluding open questions

Is it possible to use a learning dynamic to get the properties of a strong nash (learning in a cooperative sense)?