CS 6840: Algorithmic Game Theory **Spring 2017**

Lecture 7: February 8

Lecturer: Éva Tardos *Scribe: Jimmy Briggs*

7.1 Generic Price of Anarchy bounds

In this section, we consider a general game composed of the following elements.

- Players 1,2,3, $\cdots, n$.
- Strategy vector $\bar{s} = (s_1, \cdots, s_n)$, where each s_i is player i's strategy.
- $c_i(\bar{s})$, The cost to player i given that the strategies in $\bar{s}$ were chosen by their respective players.

Definition 7.1 *A vector of strategies $\bar{s}$ is a* ***Nash Equilibrium*** *if and only if $\forall i, s'$*

$$c_i(\bar{s}) \leq c_i(s'_i, \bar{s}_{-i})$$

That is to say that player i's cost only increases if she changes her strategy. Notice that this is a generic game, not necessarily a routing game. Notice also that we use the notation (s', s_{-i}) to describe player i's change to strategy s'. This notation is slightly vague, but it may help to think of s_{-i} as a copy of s with the i^{th} entry missing. i.e.

$$s_{-i} = (s_1, \cdots, s_{i-1}, -, s_{i+1}, \cdots, s_n)$$

7.2 Generalizing from routing games

Recall the Price of Anarchy proof from last time. There we assumed that our game was a routing game with linear cost functions (of the form $ax + b$). We considered two sets of paths, the Nash flow,

$$f = P_1, \cdots P_k,$$

and the optimal flow

$$f^* = P_1^*, \cdots P_k^*..$$

We also considered the cost that a player incurs in a Nash Equilibrium when switching to an optimal path. We called this $\text{cost}_{P_i^*}$.

We will use analogous terms to sketch a similar proof bounding the Price of Anarchy for general games. Instead of flows, we consider two general strategies. These are the Nash strategy,

$$\bar{s} = (s_1, \cdots, s_n),$$

and the optimal strategy,

$$\bar{s}^* = (s_1^*, \cdots, s_n^*).$$

Recall from Definition 7.1 that if $\bar{s}$ is Nash we have

$$c_i(\bar{s}) \leq c_i(s_i^*, \bar{s}_{-i})$$

Last time we used this fact, and a specific lemma to establish a Price of Anarchy bound. Let's redefine this lemma in our new, more general terms.

Lemma 7.2 *Let $\bar{s}$ and $\bar{s}^*$ be any flows in a routing game with linear edge-cost functions of the form $c_e(x) = a_e x + b_e$. Then*

$$\sum_i c_i(s_i^*, \bar{s}_{-i}) \leq \frac{5}{3} \sum_i c_i(\bar{s}^*) + \frac{1}{3} \sum_i c_i(\bar{s})$$

This still applies only to routing games, but the lemma is an example of a general property which we can use to bound PoA for any game it describes, just as we used the original form of Lemma 7.2 to bound the PoA for rounting games. We define this property now.

Definition 7.3 *A game is called **λ-μ-smooth** if there exist positive λ, μ with $\mu < 1$ such that for any strategy vectors $\bar{s}, \bar{s}^*$,*

$$\sum_i c_i(s_i^*, \bar{s}_{-i}) \leq \lambda \sum_i c_i(\bar{s}_i^*) + \mu \sum_i c_i(\bar{s})$$

As an example, we know from Lemma 7.2 that the routing game with linear costs is $\frac{5}{3}$-$\frac{1}{3}$-smooth. In words, a λ-μ-smooth game satisfies the following:

In any solution $\bar{s}$, if the solution's cost is very high compared to the optimal solution (λ cost$(\bar{s}^*)$ + μ cost$(\bar{s})$ < cost$(\bar{s})$), then there is a player i who can benefit by swapping her strategy to s_i^*.

Claim 7.4 *If a game is λ-μ-smooth, it's Price of Anarchy is bounded above by:*

$$PoA \leq \left(\frac{\lambda}{1-\mu} \right).$$

Proof: Consider a λ-μ-smooth game and let $\bar{s}$ be a Nash Equilibrium and $\bar{s}^*$ be optimal for this game. By the Nash Condition from Definition 7.1,

$$\sum_i c_i(\bar{s}) \leq \sum_i c_i(s_i^*, \bar{s}_{-i})$$

And by λ-μ-smoothness,

$$\sum_i c_i(s_i^*, \bar{s}_{-i}) \leq \lambda \sum_i c_i(\bar{s}_i^*) + \mu \sum_i c_i(\bar{s}).$$

Applying transitivity and rearranging, we see

$$(1-\mu) \sum_i c_i(\bar{s}) \leq \lambda \sum_i c_i(\bar{s}_i^*)$$

Therefore,

$$\frac{\sum_i c_i(\bar{s})}{\sum_i c_i(\bar{s}_i^*)} \leq \frac{\lambda}{1-\mu}$$

■

7.3 When does the PoA bound apply?

The proof above holds for pure Nash strategies, but what about other games? Consider, for example, the game of Rock-Paper-Scissors, a zero-sum game in which each player selects one of three options. The reward table is shown below.

	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

Rock-Paper-Scissors does not have a pure Nash Equilibrium because any deterministic strategy can be countered with another. For example, always choosing Rock loses to the strategy of alway choosing Paper. Randomized strategies, however, can achieve equilibrium.

In randomized strategies, each player i has a distribution of strategies, σ_i rather than a single strategy. Thus our solutions for randomized games are vectors of distributions rather than of strategies. That is,

$$\bar{\sigma} = (\sigma_1, \cdots, \sigma_n)$$

where each i draws independently from σ_i Now the Nash condition requires that we examine expected outcomes rather than deterministic outcomes.

Definition 7.5 *A vector of strategy distributions, $\bar{\sigma}_i$, is a **Mixed Nash Equilibrium** if and only if*

$$E_{\bar{s} \in \sigma}\left[c_i(\bar{s})\right] \leq E_{\bar{s} \in \sigma}\left[c_i(s_i', \bar{s}_{-i})\right].$$

That is to say that the expected cost for a player decreases if she switches her strategy distribution to a fixed strategy. This may seem like a strange choice of inequality. Why should we compare a mixed strategy only to pure strategies? It turns out that due to Linearity of Expectation, the above definition has the exact same effect, and is easier to work with.

7.4 Next time

Today we talked about PoA bounds in Pure Nash equilibria. But this is not the end of these bounds. Next time we will explore how to apply these bounds to Mixed Nash games and find the limit of these techniques.