CS 6840 Algorithmic Game Theory September 20, 2024

Lecture 11: Hierarchy of Equilibrium Concepts & POA Bounds in Smooth Games

Instructor: Eva Tardos *Scribe: Mohammadreza Ahmadnejadsaein*

Today's lecture covers hierarchy of equilibrium concepts and price of anarchy for learning outcomes in general finite games. We also give a general recipe for bounding price of anarchy for no-regret learning outcomes.

1 Hierarchy of Equilibrium Concepts

First let us review the definition and properties of each equilibrium concept we have learned so far. In Figure 1, we have a diagram demonstrates the hierarchy of equilibrium concepts.

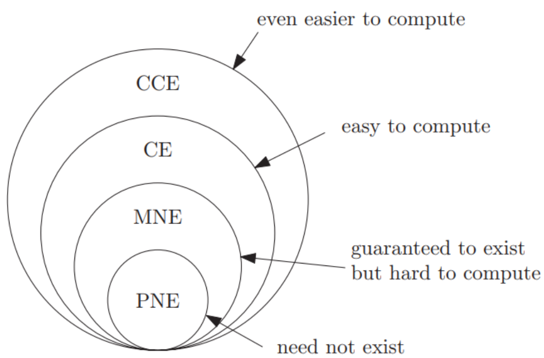

Figure 1: A hierarchy of equilibrium concepts

Pure Nash Equilibrium (PNE): A strategy profile s^* of a cost-minimization game is a pure Nash equilibrium (PNE) if for every agent $i \in \{1, \ldots, k\}$ and every unilateral deviation $s'_i \in S_i$:

$$C_i(s^*) \leq C_i(s'_i, s^*_{-i})$$

- PNE may not exist.

Mixed Nash Equilibrium (MNE): Distributions $\sigma_1, \ldots, \sigma_k$ over strategy sets $S_1, \ldots, S_k$ of a cost-minimization game constitute a mixed Nash equilibrium (MNE) if for every agent $i \in \{1, \ldots, k\}$ and every unilateral deviation $s'_i \in S_i$:

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

- Every PNE is a MNE.
- MNE always exist if the number of players and the number of strategies is finite.
- Finding MNE is known to be computationally hard, but it is not NP-complete.

Correlated Equilibrium (CE): A distributions σ on the set $S_1, \times \cdots \times S_k$ of outcomes of a cost-minimization game constitute a correlated equilibrium (CE) if for every agent $i \in \{1, \ldots, k\}$, strategy $s_i \in S_i$ and deviation $s'_i \in S_i$:

$$\mathbb{E}_{s\sim\sigma}[C_i(s) \mid s_i] \leq \mathbb{E}_{s\sim\sigma}[C_i(s'_i, s_{-i}) \mid s_i]$$

- Every MNE is a CE.
- CE is computationally tractable.
- We can interpret CE, as an equilibrium when a coordinator gives advise to each player and each players best interest is to obey the advise.

Coarse Correlated Equilibrium (CCE): A distributions σ on the set $S_1, \times \cdots \times S_k$ of outcomes of a cost-minimization game constitute a coarse correlated equilibrium (CCE) if for every agent $i \in \{1, \ldots, k\}$ and every unilateral deviation $s'_i \in S_i$:

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

- Every CE is a CCE.
- Finding a CCE is computationally tractable, however, finding the best CCE is NP-complete.
- No-regret learning algorithms can approximately find a CCE in finite steps.

2 POA in Smooth Games

In this section we give a standard recipe for bounding price of anarchy (POA).
Recall the Hotelling game where we have N locations $\{x_1, \ldots, x_N\}$, and K players each choose a location to build a Hotel there, denote s_i the location chosen by player i, and $w_x \geq 0$ as the potential maximum demand for location x, i.e., this can be number of people within maximum distance d from location x. We assume each person in demand side chooses the closest hotel within distance d form his current location.
The utility for each player would be:

$$u_i = \text{number of people choosing } s_i$$

And we can represent social welfare by:

$$SW(s) = \sum_{i=1}^{K} u_i$$

Let s be a solution and s^* be the social optimum strategy, in lecture 3, we proved the followings:

1. If s is a PNE: $u_i(s) \geq u_i(s^*_i, s_{-i})$
2. We proved that for any strategy [1] s: $\sum_{i=1}^{K} u_i(s^*, s_{-i}) \geq SW(s^*) - SW(s)$

[1]Not necessarily a Nash strategy

Combining these two inequalities gives us a lower bound for social welfare in Nash equilibrium:

$$\textit{Combine } 1. \textit{ and } 2. \Rightarrow SW(s) \geq \frac{1}{2} SW(s^*)$$

here s is a PNE, since the first inequality holds for pure Nash strategies.

2.1 POA bounds for no-regret learning algorithms

We have already discussed that the outcome of no-regret learning algorithms are approximately a coarse correlated equilibriums. Thus, a natural question arises whether there is a lower bound on POA for outcome of no-regret learning algorithms or not. In this section we will introduce a new concept to answer this question for large class of games called "smooth games".

Take $s^1, \ldots, s^t, \ldots$ outcomes of a no-regret learning algorithm in a game with K players. Following a similar approach, as we discussed in previous section leads us to these two inequalities:

1. $\sum_{t=1}^{T} u_i(s^t) \geq \sum_{t=1}^{T} u_i(s_i^*, s_{-i}^t) - reg(T)$, where $reg(T)$ is the regret of action sequence $s^1, \ldots, s^T$.
2. We proved that for any strategy s: $\sum_{i=1}^{K} u_i(s^*, s_{-i}) \geq SW(s^*) - SW(s)$

Combining these two inequalities gives us a lower bound on social welfare for outcome of learning algorithm after T iterations.

$$\begin{aligned}
\textit{Combine } 1. \textit{ and } 2. \Rightarrow \sum_{t=1}^{T} SW(s^t) &= \sum_{t=1}^{T} \sum_{i=1}^{K} u_i(s^t) \\
&\geq \sum_{i=1}^{K} \left(\sum_{t=1}^{T} u_i(s_i^*, s_{-i}^t) - reg(T) \right) \\
&= \sum_{t=1}^{T} \left(\overbrace{\sum_{i=1}^{K} u_i(s_i^*, s_{-i}^t)}^{\geq SW(s^*) - SW(s^t)} \right) - K \times reg(T) \\
&\geq \sum_{t=1}^{T} \left(SW(s^*) - SW(s^t) \right) - K \times reg(T)
\end{aligned}$$

$$\Rightarrow \boxed{\frac{1}{T} \sum_{t=1}^{T} SW(s^t) \geq \frac{1}{2} SW(s^*) - \frac{K \times reg(T)}{2T}}$$

Since in no-regret algorithms $\lim_{T \to \infty} \frac{reg(T)}{T} = 0$, for any $\epsilon > 0$ the POA for outcome of such algorithms would be greater than $\frac{1}{2} - \epsilon$ after enough number of iterations.

2.2 Smooth Games

In this section, we present a general approach for deriving bounds on the Price of Anarchy (POA) for a class of games known as "smooth games."

(λ, μ)-smooth game: a K player utility maximization game is (λ, μ)-smooth if:

$$\sum_{i=1}^{K} u_i(s_i^*, s_{-i}) \geq \lambda \cdot OPTSW - \mu \cdot SW(s)$$

where $OPTSW$ is the maximum possible social welfare, and s is any arbitrary strategy.
Theorem (POA bound of (λ, μ)-smooth game): In every (λ, μ)-smooth game utility-maximization game, the POA of CCE[2] is at least $\frac{\lambda}{1+\mu}$.
Proof. Consider a (λ, μ)-smooth utility-maximization game, a coarse correlated equilibrium σ, and an optimal outcome s^*.

$$\begin{aligned}
&\mathbb{E}_{s\sim\sigma}[SW(s)] = \mathbb{E}_{s\sim\sigma}[\sum_{i=1}^{K} u_i(s)] = \sum_{i=1}^{K} \mathbb{E}_{s\sim\sigma}[u_i(s)] \\
&\geq \sum_{i=1}^{K} \mathbb{E}_{s\sim\sigma}\left[u_i(s_i^*, s_{-i})\right] = \mathbb{E}_{s\sim\sigma}\left[\overbrace{\sum_{i=1}^{K} u_i(s_i^*, s_{-i})}^{(\lambda,\mu)-smooth}\right] \\
&\geq \mathbb{E}_{s\sim\sigma}\left[\lambda \cdot OPTSW - \mu \cdot SW(s)\right] = \lambda \cdot OPTSW - \mathbb{E}_{s\sim\sigma}\left[\mu \cdot SW(s)\right] \\
&\Rightarrow \mathbb{E}_{s\sim\sigma}\left[SW(s)\right] \geq \frac{\lambda}{1+\mu} OPTSW \iff POA \geq \frac{\lambda}{1+\mu}
\end{aligned}$$

Therefore, for any (λ, μ)-smooth game we can automatically get $\frac{\lambda}{1+\mu}$ lower bound for POA on its coarse correlated equilibrium.

3 Price of anarchy in auctions

In the last part of this session we study an example of a first price auction where K players each have values $v_1, \ldots, v_K$, and player i bids b_i.
Consider the auction winner $i^* = \arg\max\{b_1, \ldots, b_K\}$ who pays $p = b_{i^*}$. The social welfare would be:

$$SW = \overbrace{\sum_{i=1}^{K} (v_i - b_i) 1_{\{i=i^*\}}}^{\text{auction players' utility}} + \overbrace{(p)}^{\text{auctioneer's utility}} = (v_{i^*} - p) + p = v_{i^*}$$

It is easy to verify that given $v_1 > \cdots > v_K$, then $b_1 = v_2, b_i = v_i \ \forall i \geq 2$ is a pure Nash equilibrium for the game[3]. This equilibrium achieves the highest possible social welfare because as we showed $SW = v_{i^*} \leq v_1$ in this case. Therefore, in this example $POA = 1$ for the pure Nash equilibria.

[2]This is a nice property because any no-regret learning outcome converges to a CCE.
[3]Assume in the case player 1 and 2 bid the same highest amount, auctioneer always chooses player 1's bid.