Algorithmic Game Theory March 30, 2016

Lecture 5

Lecturer: Michal Feldman *Scribe: Doron Sobol, Shahar Davidson and Roee Sheffer*

1 No-Regret-Dynamics

Definition 1 (Multiplayer No-Regret-Dynamics) A no-regret game with n players. In each point in time $t = 1 \to T$:

1. Each player i (simultaneously), chooses a mixed-strategy σ_i^t, using a no-regret algorithm.
2. Each player i receives a cost vector c_i^t, where $c_i^t(s_i)$ is the expected cost of strategy s_i, with the rest of the players play according to the strategies chosen in step one. i.e. $c_i^t(s_i) = E_{s^{-i} \sim \sigma_{-i}^t}[c_i(s_i, s_{-i})]$ where $\sigma_{-i}^t = \prod_{j \neq i} \sigma_j^t$.

Claim 2 We shall denote $\sigma^t = \prod_{i=1}^{n} \sigma_i^t$, the strategy distribution of all the players at time t. Similarly, we shall denote $\sigma = \frac{1}{T}\sum_{t=1}^{T} \sigma^t$, the mean distribution over the game history. If after playing t no-regret iterations, all players have regret of at most ϵ, then σ is $\epsilon - CCE$ (ϵ-coarse-correlated equilibrium), i.e.

$$\forall s_i' \in S_i \,\Big|\, E_{s \sim \sigma}[c_i(s)] \leq E_{s_{-i} \sim \sigma_{-i}}\left[c_i(s_i', s_{-i})\right] + \epsilon$$

Proof: For each player i:

1. $E_{s \sim \sigma}[c_i(s)] = \frac{1}{T}\sum_{t=1}^{T} E_{s \sim \sigma^t}[c_i(s)]$.
2. $E_{s_{-i} \sim \sigma_{-i}}[c_i(s_i', s_{-i})] = \frac{1}{T}\sum_{t=1}^{T} E_{s_{-i} \sim \sigma_{-i}^t}[c_i(s_i', s_{-i})]$.

We will notice that equation 1 is the average over time of the expected cost of player i playing according to the no-regret algorithm. Equation 2 is the average over time of the expected value of player i playing according to strategy s_i', and the remaining players playing according to the no-regret strategy. We will notice that by definition of no-regret, the difference between the two equations is ϵ as needed. ■

2 Auctions

2.1 Single Item Auctions

Definition 3 (Single Item Auction) The single item auction model consists of a single item up for sale and n potential buyers. Formally:

- One item is being auctioned.
- $N = \{1, \ldots, n\}$ is the set of players.
- $v_i \in \mathbb{R}$ is the value, or *willingness to pay* of the ith player
- For some payment p, the utility function $u_i \in \mathbb{R}$ of the ith player is defined as

$$u_i(x,p) = \begin{cases} v_i - p & x_i = 1 \\ 0 & otherwise \end{cases}$$

  for $x = \{x_i, \ldots, x_n\} \in X$ where X is the group of feasible allocations where all values x_i are zero except for the value corresponding to the auction winner.

$$X = \left\{ x = \{x_1, \ldots, x_n\} \middle| x_i = \begin{cases} 1 & \text{player i is the auction winner} \\ 0 & otherwise \end{cases}, \sum_{i=1}^{n} x_i \le 1 \right\}$$

Definition 4 (Sealed Envelope Auction) Each player i makes a private bid b_i to the auction coordinator in a sealed envelope. Upon receiving all submitted bids, $b = \{b_1, \ldots, b_n\}$, the auction coordinator decides:

1. The allocation rule (who won the auction).
2. Payment collected from each player.

Rational players look to maximize their utility function while the auction coordinator wishes to maximize the social welfare:

$$SW = \sum_{i=1}^{n} v_i \cdot x_i$$

An auction mechanism is defined by its allocation and payment rules. In the following cases, the auction mechanism is known to all and the player values are private.

Definition 5 (First Price Auction) The winner of the auction is the highest bidder which pays the same bid.

This auction mechanism is problematic because there is no incentive for telling the truth, because if so the bidders utility will always be zero. This makes devising a strategy difficult especially in light of the fact that the other player values are supposedly private.

Definition 6 (Second Price Auction) The winner of the auction is the highest bidder which pays the second highest bid.

Definition 7 (Truthful / Incentive-Compatible Auction) An auction where telling the truth is a dominant strategy of the auction participants ($b_i = v_i$).

Claim 8 The second price auction is incentive-compatible.

Proof: Consider some player i with value v_i and the bids of the remaining players, b_{-i}. Denote $B = \max_{j \neq i} b_j$, the highest bid made by any player other than i. The player value v_i satisfies one of the following conditions:

$v_i < B$: The best response for agent i is a losing bid because bidding more than B will result in a negative utility function value. This can be attained by bidding truthfully: $b_i = v_i$.

$v_i \geq B$: Any winning bid will achieve the same non-negative utility function value, $v_i - B$, which is the highest possible. Once again, this can be attained by bidding truthfully.

■

Definition 9 (Individually Rational) An auction is individually rational if the utility function value for all rational players is non-negative. Formally,

$$\forall i, u_i = v_i - b_i \geq 0$$

.

Claim 10 Second price auctions are individually rational.

Proof: The dominating strategy of rational players in a second price auction is telling the truth. The utility of all players that do not win the auction is zero. The utility of the auction winner is $u_i = v_i - p(b) = v_i - B \geq 0$. ■

2.2 Auction Properties

Auction mechanisms can be defined by their allocation and payment rules, $M = (X(b), P(b))$ where b is the bid vector.

The following properties are desirable when planning an auction mechanism:

1. Dominant Strategy Incentive Compatible (DSIC) - the dominant strategy of rational players is truth telling.

2. Economic Efficiency - social welfare is maximized for rational players.
3. Computational Efficiency - the auction allocation rule and payment can be calculated in polynomial time.

All three properties apply to the single-item second-price auction.

Calculating the allocation rule that maximizes the social welfare can be hard. Computation efficiency is preserved by using approximation algorithms.

2.3 Search Engine Advertisement Allocation

A main source of profit for search engines is selling advertisement slots on the search engine web page. Different slots on the page are valued differently and advertisers bid in order to win advertisement slots for their ads.

Definition 11 (Advertisement Slot Auction)

- n players (advertisers) and k items (advertisement slots).
- CTR (Click Through Rate): α_j, the probability that the advertisement at slot j will be clicked. We assume without loss of generality that $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$.
- v_i, the value per click of the player i. This value is private.
- The value given by player i to occupying slot j is $\alpha_j \cdot v_i$.

Social welfare can be maximized by sorting the player values $v_1 \geq v_2 \geq \cdots \geq v_k$ and allocating slot j to the player with value v_j (after sorting). In this case, we shall achieve the social welfare value:

$$SW = \sum_{j=1}^{k} \alpha_j v_j$$

The specified algorithm maximizes the social welfare and allocation can be calculated efficiently. Next lecture we will attempt to find a payment rule for which the rest of the desirable auction properties hold.

2.4 Single Parameter Environment

The single parameter environment consist of n players, where each player has a value v_i per allocated item. Item allocation is designated by $x = \{x_1, \ldots, x_n\} \in X$ where every x_i is the item allocation count of player i and X is the group of feasible allocations.

In the case of a **single parameter auction**:

- Input: $b = \{b_1, \ldots, b_n\}$, the player bids.
- Output: $x(b) \in X \subseteq \mathbb{R}^n$ and $p(b) \in \mathbb{R}^n$, the allocation and payment vectors.
- The utility function of player i: $u_i(b) = v_i \cdot x_i(b) - p_i(b)$. In this course, we shall assume that $p_i(b) \in [0, b_i \cdot x_i(b)]$.

Definition 12 An allocation rule x is **implementable** if a payment rule p exists such that the mechanism $M(x, p)$ is DSIC.

Definition 13 An allocation rule x is **monotone** if for every player i and bid vector b_{-i}, the allocation $x_i(z, b_{-i})$ is non-decreasing in z.