Lecture 1

Lecturer: Michal Feldman *Scribe: Yuval Rochman, Amir Rubinstein, Omri Sharabi*

1 Administration

- Lecturer: Michal Feldman- mfeldman@tau.ac.il
- TA: Shai Vardi- shaivar1@tau.ac.il
- Reception hour: Sunday, 12 : 30, Schreiber 121
- Additional references, problem set assignments and other information can be found at the course website: http://agttau.wikidot.com/.

Today's topics:

- Introduction
- Formal definitions in game theory
- Nash equilibrium
- Zero-sum games

2 Course Requirements

- **Scribes:** Every student must write a lecture scribe. Each scribe should be written by no more than three students.
- **Problem set assignments:** Submitting homework assignments is mandatory. Students may work in groups, the partners' names must be written on the assignment. Assignments are submitted to Shai Vardi's cell on the second floor in Schreiber.
- **Final Project:** Can be either a survey of articles, or a research project. The project can be done in groups, and must contain no more than 10 one-sided pages.

3 Introduction

Algorithmic Game Theory has emerged in the last years as a new developing field done at the intersection of Computer Science, Game Theory and Economic Theory. The field was motivated by the emergence of the internet. In this course we will discuss the following objectives:

1. Mechanism Design - How to design the game mechanism in order to achieve a desired goal. We assume that users are *rational*, i.e., players maximize their utility.
2. The inefficiency of a game due to selfish behavior of the rational players.
3. The computation complexity of the proposed algorithms, and compute efficient solutions.

3.1 Mechanism Design

The first part of the course analyzes the mechanism design of games in which rational players participate.

To begin, let us examine the problem of allocating a single item among various buyers. There are several plausible auction formats even for ther simple problem. The goal of each format is to allocate the item to the buyer that values it the most. The formats differ in their payment rules (e.g, first- or second-price), and the actual payment of the bidders.

For every bidder i we denote by v_i and b_i the bidder's value and the bidder's bid for that item. Let p_i be the payment of the i^{th} bidder, which is less than or equal to b_i if the bidder wins and otherwise 0. We denote by u_i the utility of bidder i which is

$$u_i = \begin{cases} v_i - p_i & \text{if bidder} i \text{ wins} \\ o & \text{otherwise.} \end{cases} \tag{1}$$

For the single-item auction we define two different methods of payment:

1. **First Price Auction-** In a First Price Auction the bidder with the maximum bid wins the item, and pays the amount of her bid to the seller. That means, the winner is bidder $i_1 = \arg\max_i b_i$ and pays $p_{i_1} = b_{i_1}$.
2. **Second Price Auction (Vickrey auction)-** In a Second Price Auction the bidder with the highest bid wins the item, but the price paid is the second-highest bid. That means bidder $i_1 = \arg\max_i b_i$ wins and pays $p_{i_1} = b_{i_2}$, where bidder $i_2 = \arg\max_{i \neq i_1} b_i$ gave the second-highest bid.

In both auctions, if there are more than two bidders submitting the maximum bid, then the winner is chosen at random.

The main problem with a First Price Auction is that a rational bidder might not bid her item valuation ($b_i \neq v_i$). For example, if the valuations of all bidders are common knowledge, then the bidder with the maximum valuation would state the second-highest valuation as her bid plus ϵ. That bidder would win the auction with a higher utility than she would bid her true valuation.

On the other hand, in a Second Price Auction the bidders best strategy is to state their own valuations, which is proven by the following theorem.

Theorem 1 *In a Second Price Auction each rational bidder i maximizes her utility by bidding (revealing) her true valuation, i.e., $b_i = v_i$.*

Proof: There can be three cases: 1) $v_i > \max_{j \neq i} b_i$, 2) $v_i < \max_{j \neq i} b_i$ 3) $v_i = \max_{j \neq i} b_i$. We will prove that all of in these cases the optimal strategy is to bid on the true valuation.

If $v_i > \max_{j \neq i} b_i$ then bidder i can win with a constant positive utility of $u_i = v_i - \max_{j \neq i} b_i > 0$ or lose with 0 utility. If bidder i bids on her true valuation ($b_i = v_i$) then she wins and her utility is maximized.

If $v_i < \max_{j \neq i} b_i$ then bidder i can win with a constant negative utility of $u_i = v_i - \max_{j \neq i} b_i < 0$ or lose with 0 utility. If $b_i = v_i$ then the bidder loses and her utility is maximized.

If $v_i = \max_{j \neq i} b_i$, then bidder i has 0 utility, whether the bidder loses or wins. Thus, bidding the true valuation maximizes the bidder utility. ■

3.2 Inefficiency due to the selfishness of players

The second part of the course measures how the efficiency of a system degrades due to selfish behavior of its player, as shown in the following example:

Example 1 *Consider the four-node network shown in Figure 1. There are two disjoint routes from s to t , each with a combined cost of $x+1$, where x is the amount of traffic that uses the route. Assume that there is one unit of traffic. In the equilibrium flow, the traffic is split evenly between the two routes, and all of the traffic experiences $3/2$ units of cost. Now suppose that, in an effort to decrease the cost encountered by the traffic, we build a zero-cost edge connecting w to v. The new network is shown in Figure 2. The previous equilibrium flow does not persist in the new network: the cost of the new route $s \to w \to v \to t$, which equals to $2 \cdot x$, is never worse than that along the two original paths (as $2 \cdot x \leq x + 1$). As a consequence, the unique equilibrium flow routes all of the traffic on the new route. Because of the ensuing heavy congestion on the edges (s, v) and (w, t), all of the traffic now experiences two units of cost. Braesss Paradox thus shows that the intuitively helpful action of adding a new zero-cost edge can increase the cost experienced by all of the traffic!*

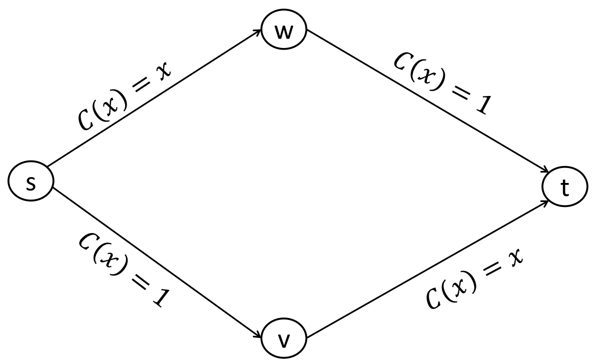

Figure 1: Braess's paradox: The stable and efficient state

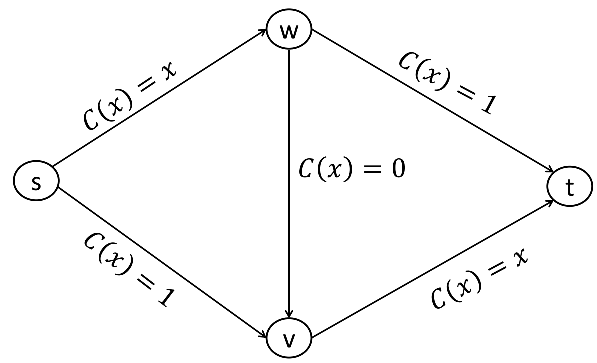

Figure 2: Braess's paradox: The inefficient state

For this example, we define the *Price of Anarchy* (PoA) as the cost experienced as a result of the players selfishness, divided by the cost in the optimal case, which is

$$PoA = \frac{2}{1.5} = 4/3$$

In the course we will prove that no non-atomic instance with linear cost functions has a price of anarchy larger than 4/3.

4 Mechanism design - modeling

For the study and analysis of design mechanism we use the following definitions:

Definition 2 *Let S_i be the strategy set of player i and u_i the utility of player i. A* ***Game*** *of n players is a pair (S, u) where $S = S_1 \times S_2 \times \cdots \times S_n$ is the set of strategy profiles and $u(s) = (u_1(s), u_2(2), \ldots, u_n(s))$ is the utility function. We denote by S_{-i} the strategy profiles of all players except player i. The strategy profile $s = (s_{-i}, s_i)$ is composed of the s_i strategy of player i, and a strategy profile $s_{-i} \in S_{-i}$ of other players.*

Definition 3 *Let $x, x^* \in S_i$ be two strategies of player i. Then x* dominates *x^* if choosing x always gives as good as or a better outcome than choosing x^*. Formally, for every strategy profile $s_{-i} \in S_{-i}$ we have $u_i(x, s_{-i}) \geq u_i(x^*, s_{-i})$*

Definition 4 *A strategy $s_i \in S_i$ is called a* ***dominating strategy*** *for player i if it dominates all other strategies of player i.*

Definition 5 *A strategy $x \in S_i$ is a **Best Response** to a strategy profile $s_{-i} \in S_{-i}$ if $\forall x^* \in S_i : u_i(x, s_{-i}) \geq u_i(x^*, s_{-i})$*

Definition 6 *A game with two players is a **Zero-Sum Game** if for every strategy $s \in S$ the sum of the players' utilities is zero, i.e. $u_1(s) + u_2(s) = 0$.*

5 Nash Equilibrium in pure strategies

Definition 7 *The profile $s = (s_1, s_2, \ldots, s_n)$ is in Nash equilibrium (NE) if for every player i, s_i is a best response to s_{-i}.*

In Nash Equilibrium no player can improve their utility by deviating to another strategy, assuming that all other players do not change their own strategies. We note that the notion of dominating strategies equilibrium is more strict: if there is an equilibrium with dominating strategies, this is also a NE.

For example, in the prisoner's dilemma, (D,D) is a Nash equilibrium (but clearly not the optimal profile of strategies, as (C,C) delivers a better result for both players. However (C,C) is not a NE as both players have incentives to change their strategy). As can be seen, this game has only one NE. The "Battle of Sexes" game has two Nash equilibria: (O,O) and (B,B), while Rock-Paper-Scissors has none. The latter can be generalized:

Claim 8 *In a zero-sum game there can be no pure NE.*

6 Mixed Strategies and Nash theorem

Denote $\Delta(s) = \{x_1, \ldots, x_n \mid \forall j \ x_j \geq 0, \ \sum_{j=1}^{n} x_j = 1\}$. A mixed strategy of player i is $x^i = (x^i_1, \ldots, x^i_{|S_i|}) \in \Delta(s_i)$, where x^i_j is the probability that player i plays strategy j.

The expected payoff of player i is $u_i(x^1, \ldots, x^n) = E_{s_j \sim x^j}[u_i(s_1, \ldots, s_n)]$. A rational player in a game with mixed strategy seeks to maximized their expected payoff.

Definition 9 *A NE in mixed strategies is $(x^1, \ldots, x^n)$ s.t. every x^i is a best response to x^{-i}.*

Note that if for every player i there is some j s.t. $x^i_j = 1$, this is the private case of pure NE.

Is there a mixed strategy NE in the Rock-Paper-Scissors game? The mixed strategy $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for both players is such an equilibrium, since each player earns 0, and has no incentive to change these probabilities. Note that if there is a mixed strategy NE, each of the strategies with positive probabilities must have the same payoff (otherwise a deviation would be profitable).

Theorem 10 (Nash's Existence Theorem ('51)) *In a mixed strategies game with a finite number of players and a finite number of strategies there is at least one mixed NE.*

The proof of Nash's theorem is based on Brouwers Fixed-Point theorem (which we will not prove):

Theorem 11 *Brouwer Fixed-Point Theorem: Let C be a closed bounded convex set in $\mathbb{R}^t$ and $f : C \to C$ a continuous function, then there exists a point $x \in C$ s.t. $f(x) = x$.*

Example: $C = [0, 1]$, and $f : [0, 1] \to [0, 1]$ some continuous function.
Define $g(x) = f(x)-x$. Then $g(0) = f(0) \geq 0$ and $g(1) = f(1)-1 \leq 0$. By the intermediate value theorem there exists x s.t. $g(0) = 0$, therefore $f(x) = x$.

Proof: (Nash Existence Theorem)

We are looking for a profile of mixed strategies $(x^1, \ldots, x^n)$ s.t. every x^i is a best response to x^{-i}. We may want to consider $(x^1, \ldots, x^n) \to (BR(x^{-1}), \ldots, BR(x^{-n}))$ to apply the Fixed-Point Theorem (concluding that some profile is mapped to itself, thus is a NE). However, this is not necessarily a funtion (there may be several best responses of a player to the other players' strategies). Therefore, we will denote:

$$c_i^j = u_i(j, x^{-i}) - u_i(x^i, x^{-i})$$

This is the value player i earns from deviating to the (pure) strategy j. We limit this notion to profitable deviations only:

$$c_i^{j+} = max(c_i^j, 0)$$

Finally, we can define the following continuous function, to which the Brouwer Fixed-Point Theorem is applicable:

$$\hat{x}_i^j = \frac{x_i^j + c_i^{j^+}}{1 + \sum_{j=1}^{|S_i|} c_i^{j^+}}$$

(the denominator normalizes this expression to be a proper distribution).

We conclude that $x_i^j = \hat{x}_i^j$ for the fixed point x. It is left to show that in the fixed point, $\forall i, j \;\; C_i^{j+} = 0$. Assume by contradiction the existence of i, j s.t. $C_i^{j+} > 0$. $\forall i \;\; \sum_{j=1}^{|S_i|} c_i^j x_i^j = 0$. By definition $\hat{x}_i^j > 0$, and so $x_i^j = \hat{x}_i^j > 0$. Therefore, there must exist j' s.t. $C_i^{j'} < 0$ and $x_i^{j'} > 0$, and by definition $C_i^{j'+} = 0$ for that j. This leads to the following:

$$x_i^{j'} = \hat{x}_i^{j'} = \frac{x_i^{j'} + C_i^{j'+}}{> 1} < x_i^{j'},$$

which is a contradiction. ■

Comment: Computationally, finding NE is PPAD-complete (a class believed to be in NP/P).

7 Two person Zero-Sum game

In zero-sum games the gain of one player is the exact loss of the other player. Therefore, in a matrix game there is no need to write the payoff of both players in the cells. It is enough to write the gain of the first player as we know that the first player aims at maximizing this value and the second player aims at minimizing the value, we often refer to the first player as the maximizer and to the second player as the minimizer. As an example, onsider the following zero-sum game:

	L	M	R
U	0	20	100
D	10	1	200

The row player is the maximizer and has two pure strategies, U and D. The column player is the minimizer and has three pure strategies, L, M and R.
$a_{i,j}$ is the row player's payoff (column player's loss) if the row player chooses strategy i and the column player chooses strategy j

In the case that the players take turns, if the maximizer plays first she will choose the row that offers her the maximal minimum payoff; therefore she will choose D. The minimizer will then choose M and the result of the game will be 1. If the minimizer plays first she will choose the column that offers her the minimal maximum loss possible, therefore she will choose L, the maximizer will then choose D and the result of the game will be 10.

More formally, if the minimizer plays first she will choose the row: $r = \arg\min_i \max_j a_{i,j}$, if the maximizer play first it will choose the column: $s = \arg\max_j \min_i a_{i,j}$

Lemma 12 $\max_j \min_i a_{i,j} \leq \min_i \max_j a_{i,j}$

Proof: It's enough to show that for all $\hat{\imath}, \hat{\jmath}$ it holds that $\min_j a_{\hat{\imath},j} \leq \max_i a_{i,\hat{\jmath}}$. this is true because $\min_j a_{\hat{\imath},j} \leq a_{\hat{\imath},\hat{\jmath}} \leq \max_i a_{i,\hat{\jmath}}$ ■

8 John von Neumann MinMax theorem for mixed strategies

- Denote the mixed strategies of player 1: $X = (x_1, x_2, \ldots, x_n)$
- Denote the mixed strategies of player 2: $Y = (y_1, y_2, \ldots, y_m)$
- The gain of player 1 (the loss of player 2) (i.e, the game value) is: $v_1 = \sum_i \sum_j x_i y_j a_{i,j}$

If player 2 knows the mixed strategy of player 1 (X) she will choose, without loss of generality, a pure strategy that minimizes $\sum_i x_i a_{i,j}$.
The mixed strategy of player 1 if she plays first is: $X = arg \max_s \min_j \sum_i s_i a_{i,j}$
Symmetrically, if player 1 know the mixed strategy of player 2 (Y), she will choose, without loss of generality, a pure strategy that minimizes $\sum_j y_i a_{i,j}$.
The mixed strategy of player 2 if she plays first is: $Y = arg \min_r \max_i \sum_j r_j a_{i,j}$

Theorem 13 *John von Neumann MinMax theorem for mixed strategies: the game value is* $v = \max_x \min_j \sum_i x_i a_{i,j} = \min_y \max_i \sum_j y_j a_{i,j}$

(we will prove this theorem in the next lecture). It follows that (X, Y) is a Nash equilibrium for the mixed strategies X and Y, meaning each player has a value it can guarantee regardless of what the other player does. In a zero sum two player game, if we find the game value we can find a Nash equilibrium.