Computational Learning Theory	Spring Semester, 2005/6

Lecture 1: March 7

Lecturer: Yishay Mansour *Scribe: Yuval Netzer, Shahar Fattal*

1.1 Introduction

Several fields in computer science and economics are focused on the analysis of Game theory. Usually they observe Game Theory as a way to solve optimization problems in systems where the participants act independently and their decisions affect the whole system.
Following is a list of research fields that utilize Game Theory:

- Artificial Intelligence (AI) - There has been a shift from interest in Machine Learning, where one builds a hypothesis based on data, to Reinforcement Learning, which deals with a single agent (robot), to Multi Agent learning.
- Communication Networks - The goal has been to suggest mechanisms, where it will be both in the user's interest to follow them, and the outcome will be efficient.
- Computer Science Theory - There are several subfields that use Game Theory:
  - Algorithm Mechanism Design.
  - Complexity Issues.
  - Quality of Nash Equilibrium (compared to optimal solution).

1.2 Course Syllabus

1. Introduction
2. Quality of an Equilibrium
   (a) Job Scheduling (Price of Anarchy)
   (b) Routing (Price of Anarchy)
   (c) Network Creation (Price of Anarchy)
   (d) Network Design (Price of Stability)
3. Equilibrium Existence

   (a) Two players zero sum game
   (b) Correlated equilibrium
   (c) Congestion and potential games.
   (d) Nash Equilibrium (existence)
   (e) Graphical games and hardness

4. Repeated Games
   (a) General folklore theorems
   (b) Internal and External Regret
   (c) Option pricing
   (d) Dynamics of reaching equilibrium
   (e) Vector payoff games (Approachability Theorem)

5. Mechanism Design
   (a) Social choice
   (b) Maximizing social welfare(VCG mechanism)
   (c) Auctions (combinatorial)
   (d) Auctions (digital goods)

1.3 Strategic Games

A strategic game is a model for decision making where there are N players, each choosing an action. A player's action is chosen once and cannot be changed afterwards.

Each player i can choose an action a_i from a set of actions A_i. let A be the set of all possible action vectors $\times_{j \in N} A_j$. Thus, the outcome of the game is an action vector $\vec{a} \in A$.

All the possible outcomes of the game are known to all the players and each player i has a preference relation over the different outcomes of the game: $\vec{a} \preceq_i \vec{b}$ for every $\vec{a}, \vec{b} \in A$. The relation $\vec{a} \preceq_i \vec{b}$ implies that player i prefers $\vec{b}$ over $\vec{a}$ (or has equal preference for either).

Definition A **Strategic Game** is a triplet $\langle N, (A_i), (\preceq_i) \rangle$ where N is the number of players, A_i is the finite set of actions for player i and $\preceq_i$ is the preference relation of player i.

We will use a slightly different notation for a strategic game, replacing the preference relation with a payoff function $u_i : A \to \mathbb{R}$, where $\vec{a} \preceq_i \vec{b} \iff u_i(\vec{a}) \leq u_i(\vec{b})$. The player's target is to maximize her own payoff. Such strategic game will be defined as: $\langle N, (A_i), (u_i) \rangle$.

This game theoretic model is very abstract. Players can be humans, companies, governments etc. The preference relation can be subjective evolutional etc. The actions can be simple, such as "go forward" or "go backwards", or can be complex, such as design instructions for a building.

Several player behaviors are assumed in a strategic game:

- The game is played only once.
- Each player "knows" the game (each player knows all the actions and the possible outcomes of the game).
- The players are rational. A rational player is a player that plays selfishly, wanting to maximize her own benefit of the game (the payoff function).
- All the players choose their actions simultaneously (but do not know the other players current choices).

1.4 Pareto Optimality

An outcome $\vec{a} \in A$ of a game $\langle N, (A_i), (u_i) \rangle$ is Pareto Optimal if there is no other outcome $\vec{b} \in A$ that makes every player at least as well off and at least one player strictly better off. That is, a Pareto Optimal outcome cannot be improved upon without hurting at least one player.

Definition An outcome $\vec{a}$ is **Pareto Optimal** if there is **no** outcome $\vec{b}$ such that
$\forall_{i \in N}\ u_i(\vec{a}) \leq u_i(\vec{b})$ and $\exists_{j \in N}\ u_j(\vec{a}) < u_j(\vec{b})$.

1.5 Nash Equilibrium

A Nash Equilibrium is a state of the game where no player prefers a different action if the current actions of the other players are fixed.

Definition An outcome a^* of a game $\langle N, (A_i), (\preceq_i) \rangle$ is a **Nash Equilibrium** if:
$\forall_{i \in N} \forall_{b_i \in A_i}\ (a^*_{-i}, b_i) \preceq (a^*_{-i}, a^*_i)$,
where $(a_{-i}, x) = (a_1, a_2, ..., a_{i-1}, x, a_{i+1}, ..., a_n)$.

We can look at a Nash Equilibrium as the best action that each player can play based on the given set of actions of the other players. Each player cannot profit from changing her action, and because the players are rational, this is a "steady state".

Definition Player i **Best Response** for a given set of other players actions $a_{-i} \in A_{-i}$ is the set: $BR_i(a_{-i}) := \{b \in A_i | \forall_{c \in A_i} (a_{-i}, c) \preceq_i (a_{-i}, b)\}$.

Under this notation, an outcome a^* is a Nash Equilibrium if $\forall_{i \in N}\ a^*_i \in BR_i(a^*_{-i})$. We note that a Nash Equilibrium is not always Pareto Optimal, and there can be multiple Nash Equilibria.

1.6 Matrix Representation

A two player strategic game can be represented by a matrix whose rows are the possible actions of player 1 and the columns are the possible actions of player 2. Every entry in the matrix is a specific outcome and contains a vector of the payoff value of each player for that outcome. For example, if $A_1 = \{r_1, r_2\}$ and $A_2 = \{c_1, c_2\}$ the matrix representation is:

	c_1	c_2
r_1	(w_1, w_2)	(x_1, x_2)
r_2	(y_1, y_2)	(z_1, z_2)

Where for $i \in \{1, 2\}$, $w_i = u_i(r1, c1)$, $x_i = u_i(r_1, c_2)$, $y_i = u_i(r_2, c_1)$ and $z_i = u_i(r_2, c_2)$.

1.7 Strategic Games: Examples

The following are examples of two players games with two possible actions per player. The set of deterministic Nash Equilibrium points is described in each example.

1.7.1 Matching Pennies

The goal of this game is to select a winner. Both players select Head or Tails. The row player wins if they match, and the column player wins if they mismatch.

	Head	Tail
Head	$(1, -1)$	$(-1, 1)$
Tail	$(-1, 1)$	$(1, -1)$

In this game there is no Deterministic Nash Equilibrium point. Also, this is a zero sum game (the sum of the profits of each player over all possible outcomes is 0).

1.7.2 Battle of the Sexes

In this game, the two players need to coordinate on an event (sports or opera). They both prefer to go to the same event together, but they have a different preference between the events.

	Sports	Opera
Sports	$(2, 1)$	$(0, 0)$
Opera	$(0, 0)$	$(1, 2)$

There are two Nash Equilibrium points: (Sports, Sports) and (Opera, Opera). Both are also Pareto Optimal.

1.7.3 A Coordination Game

In this game, generals need to decide whether to attack or retreat. A failure to reach the same decision would result in losing the battle. If they both attack, they win. If they both retreat they have no gain or loss.

	Attack	Retreat
Attack	$(10, 10)$	$(-10, -10)$
Retreat	$(-10, -10)$	$(0, 0)$

There are two Nash Equilibrium outcomes: (Attack, Attack) and (Retreat, Retreat), where (Attack, Attack) is also Pareto Optimal.
A question that rises from this game and its equilibria is how the two players can move from one Equilibrium point, (Retreat, Retreat), to the better one (Attack, Attack). Another way to look at it is how the players can coordinate to choose the preferred equilibrium point, in this symmetric game.

1.7.4 The Prisoner's Dilemma

There are two prisoners that committed a crime. If they both do not confess, they get a low punishment. If they both confess, they get a more severe punishment. If one confesses and the other does not, then the one that confesses gets a low punishment and the other gets a very severe punishment.

	Don't Confess	Confess
Don't Confess	$(-1,-1)$	$(-4,0)$
Confess	$(0,-4)$	$(-3,-3)$

There is one Nash Equilibrium point: (Confess, Confess), which is not Pareto Optimal. Though it looks natural that the two players will cooperate, the cooperation point (Don't Confess, Don't Confess) is not a steady state since once in that state, it is more profitable for each player to move into 'Confess' action, assuming the other player will not change its action.

Definition a_i is a **Weak Dominant Strategy** for player i if

$$\forall b_{-i} \in A_{-i}.\forall b_i \in A_i : u_i(b_{-i}, b_i) \leq u_i(b_{-i}, a_i)$$

a_i is a **Strong Dominant Strategy** for player i if

$$\forall b_{-i} \in A_{-i}.\forall b_i \in A_i : u_i(b_{-i}, b_i) < u_i(b_{-i}, a_i)$$

In The Prisoner's Dilemma game the action 'Confess' is a strong dominant strategy for both players.

1.7.5 Dove-Hawk

The two players need to decide whether to engage in war with each other. If one engages in war (Hawk) and the other does not, then he has a significant gain. However, if they both engage in war, they lose.

	Dove	Hawk
Dove	$(3,3)$	$(1,4)$
Hawk	$(4,1)$	$(0,0)$

There are two Nash Equilibrium points: (Dove, Hawk) and (Hawk, Dove).

1.7.6 Rock Paper Scissors

This is a zero sum game.
There is no Nash Equilibrium point here.

	R	P	S
R	$(0,0)$	$(-1,1)$	$(1,-1)$
P	$(1,-1)$	$(0,0)$	$(-1,1)$
S	$(-1,1)$	$(1,-1)$	$(0,0)$

1.7.7 Auction

There are N players, each one wants to buy an object.

- Player i's valuation of the object is v_i, and, without loss of generality, $v_1 > v_2 > ... > v_n > 0$.
- The players simultaneously submit bids - $k_i \in [0, \infty)$. The player who submit the highest bid - k_i wins.
- In a *first price auction* the payment of the winner is the price that he bids. That is, the payoff of player i is $u_i = \begin{cases} v_i - k_i, \ i = argmax\ k_i \\ 0, \ otherwise \end{cases}$.

  One Nash equilibrium point is $k_1 = v_2 + \epsilon, k_2 = v_2, ..., k_n = v_n$. In fact one can see that $k_3, \ldots, k_n$ have no influence. Unfortunately, the first player needs to "know" the second highest bid, which in practice is not available to him.

In a *second price auction* the payment of the winner is the highest bid among those submitted by the players who do not win. Player i's payoff when he bids v_i is at least as high as his payoff when he submits any other bid, regardless of the other players' actions. Therefore, for each player i, the bid $k_i = v_i$ is a weakly dominant strategy. This suggests that it is of best interest of all players to bid their own value. Also, when all players bid their own value, it is a Nash Equilibrium. This strategy causes the player to bid truthfully.

For example, consider the two players case. We show that $\forall k_1, k_2 . u_1(v_1, k_2) \geq u_1(k_1, k_2)$.

If $v_1 > k_2$, then $u_1(v_1, k_2) = v_1 - k_2 \geq u_1(k_1, k_2) = \begin{cases} v_1 - k_2, \ k_1 > k_2 \\ 0, \ k_1 < k_2 \end{cases}$

If $v_1 < k_2$, then $u_1(v_1, k_2) = 0 \geq u_1(k_1, k_2) = \begin{cases} v_1 - k_2 (< 0), \ k_1 > k_2 \\ 0, \ k_1 < k_2 \end{cases}$

1.7.8 A War of Attrition

Two players are involved in a dispute over an object.

- The value of the object to player i is $v_i > 0$.

- There is also time $t \in [0, \infty)$.
- Each player chooses when to concede the object to the other player
- If the first player to concede does so at time t, his payoff $u_i = -t$, and the other player obtains the object at that time and his payoff is $u_j = v_j - t$.
- If both players concede simultaneously, the object is split equally, player i receiving a payoff of $\frac{v_i}{2} - t$.

The only Nash equilibria are when one of the players concede immediately and the other wins.

1.7.9 Location Game

- Each of n people chooses whether or not to become a political candidate, and if so which position to take.
- The distribution of favorite positions is given by the density function f on $[0, 1]$.
- A candidate attracts the votes of the citizens whose favorite positions are closer to his position.
- If k candidates choose the same position then each receives the fraction $\frac{1}{k}$ of the votes that the position attracts.
- Each person prefers to be the unique winning candidate rather than to tie for first place, prefers to tie the first place rather than to stay out of the competition, and prefers to stay out of the competition rather than to enter and lose.

When $n = 2$ there is a Nash Equilibrium when one of the players is in the middle, and the other player is as close as possible to him (from one of his sides).

When $n = 3$ there is no Nash equilibrium. No player wants to be in the middle, since the other players will be as close as possible to the middle player, either from the left or the right.

For a circle there is always a Nash Equilibrium, where the players are placed at equal distances.

1.8 Mixed Strategy

Now we will expand our game and let the players' choices to be stochastic. Each player $i \in N$ will choose a probability distribution P_i over A_i:

1. $P = \langle P_1, ...P_N \rangle$

2. $P(\vec{a}) = \prod P_i(a_i)$ (the choices of the players are statistically independent).

3. $u_i(P) = E_{\vec{a} \sim P}[u_i(\vec{a})]$

Note that the function u_i is linear in P_i : $u_i(P_{-i}, \lambda\alpha_i + (1-\lambda)\beta_i) = \lambda u_i(P_{-i}, \alpha_i) + (1-\lambda)u_i(P_{-i}, \beta_i)$.

Definition $\mathit{support}(P_i) = \{a | P_i(a) > 0\}$

Note that the set of Nash equilibria of a strategic game is a subset of its set of mixed strategy Nash equilibria.

Lemma 1.1 *Let* $G = \langle N, (A_i), (u_i) \rangle$. *Then* α^* *is Nash equilibria of* G *if and only if* $\forall_{i \in N} : \mathit{support}(P_i) \subseteq BR_i(\alpha^*_{-i})$, *where* $BR_i(P_{-i}) := \{P \in D(A_i) | \ \forall_{p' \in D(A_i)} \ u_i(P_{-i}, P) \geq u_i(P_{-i}, P')\}$.

Proof:
$\Rightarrow$ Let α^* be a mixed strategy Nash equilibria ($\alpha^* = (P_1, ..., P_N)$). Suppose $\exists a_i \in \mathit{support}(P_i)$ such that $a_i \notin BR_i(\alpha^*_{-i})$. Then player i can increase her payoff by transferring probability to $a'_i \in BR_i(\alpha^*_{-i})$; hence α^* is not mixed strategy Nash equilibria - contradiction.
$\Leftarrow$ Let q_i be a probability distribution s.t. $u_i(Q) > u_i(P)$ in response to a mixed strategy α^*_{-i}. Then by the linearity of u_i, $\exists b \in \mathit{support}(Q_i), c \in \mathit{support}(P_i)$. $u_i(\alpha^*_{-i}, b) > U_i(\alpha^*_{-i}, c)$; hence $c \notin BR_i(\alpha^*_{-i})$ - contradiction. □

Later on we will show that every game has a mixed strategy Nash Equilibria.

1.8.1 Battle of the Sexes (example)

As we mentioned above, this game has two deterministic Nash equilibria, (S,S) and (O,O). Suppose α^* is a stochastic Nash equilibrium:

- $\alpha^*_1(S) = 0$ or $\alpha^*_1(S) = 1 \Rightarrow$ same as the deterministic case.

- $0 < \alpha^*_1(S) < 1 \Rightarrow$ by the lemma above $2\alpha^*_2(O) = \alpha^*_2(S)$ ($\alpha^*_2(O) + \alpha^*_2(S) = 1$) and thus $\alpha^*_2(O) = \frac{1}{3}$, $\alpha^*_2(S) = \frac{2}{3}$. Since $0 < \alpha^*_2(S) < 1$ it follows from the same result that $2\alpha^*_1(S) = \alpha^*_1(O)$ so $\alpha^*_1(S) = \frac{1}{3}$, $\alpha^*_1(O) = \frac{2}{3}$.

The mixed strategy Nash Equilibrium is $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$.

1.9 Correlated Equilibrium

We can think of a traffic light that correlates, advises the cars what to do. The players observe an object that advises each player of her action. A player can either accept the advice or choose a different action. If the best action is to obey the advisor, the advice is a correlated equilibrium formally.

Definition Let Q be a probability distribution over A. Q is a Nash correlated equilibrium if $\forall i. \forall z_i \in suppport(Q). \forall x \in A_i. E_Q[U_i(a_{-i}, z_i)|a_i = z_i] \geq E_Q[U_i(a_{-i}, x)|a_i = z_i]$