Topics on the Border of Economics and Computation	October 30, 2005
Lecture 1	
Lecturer: Noam Nisan	*Scribe: Michael Schapira*

1 Introduction to Game Theory; Zero-Sum Games

The field of *decision theory* deals with problems in which there is a single person who needs to choose a course of action that would best serve his purposes out of a set of possible ways of action. In a *game*, there are at least two people. As in decision theory, the players have different courses of actions, called *strategies*, from which they can choose. However, in a game the benefit of a player does not only depend on the choice he himself makes but also on the choices made by the other players. Hence, a large part of the difficulty of games depends on the interaction between the different parties. *Game theory* is the mathematical and economic field that studies this type of conflicts.

Consider a game with two players, we shall refer to as 1 and 2. We assume that every player has preferences over the different possible outcomes of the game. In some cases these preferences are *ordinal* - that is, merely state for every two outcomes which one is preferred over the other. We shall assume that every player has *cardinal* preferences. I.e., every player assigns a numerical value (in $\Re$) to each possible outcome. We do not seek connections between the values a player assigns to the different outcomes and some "reality" but merely regard these values as representing that player. We shall refer to a player's value for an outcome as his *utility* in that outcome. The basic assumption of game theory is *rationality*: every player is only interested in maximizing their utility.

Every player $k \in \{1, 2\}$ has a set of strategies $S_i = \{s_1^k, ..., s_{m_k}^k\}$ from which he must choose one. In this lecture we will focus on zero-sum games. A zero-sum game is a game in which there is absolutely no cooperation between the players: what one player gains the other player must lose. This game can be described by a matrix $A = \{a_{ij}\}_{i \in [m_1], j \in [m_2]}$ such that a_{ij} denotes the value player 1 assigns to the outcome of the game achieved when he chooses s_i^1 and player 2 chooses s_j^2. Since we are dealing with zero-sum games this means that the value of player 2 for this outcome must be $(-a_{i,j})$. This way of representing games is called the *strategic form.* In the following lectures we shall encounter another way for the representation of games, namely the *extensive form.*

Let us look at some examples:

Example 1 *The rock-paper-scissors game*

	rock	paper	scissors
rock	0	-1	1
paper	1	0	-1
scissors	-1	1	0

The rows of the matrix represent the different strategies of the player 1 and the columns represent the strategies of player 2. Each cell's value states whether player 1 wins (1), player 2 wins (-1), or we have a tie (0). For instance, $a_{12} = -1$ since if player 1 chooses "rock" and player 2 chooses "paper" then player 2 wins.

Example 2 *Consider a chess game. Every strategy of player 1 (the white) would consist of an opening move, a move for every possible move made by player 2, a move for every second possible move chosen by player 2, and so on. Player 2's strategies are defined similarly. It is possible to represent chess in strategic form using a gigantic matrix. For every strategy of players 1 and 2 we can determine if one of them wins or if there is a tie. The values of the cells would then be 1 (player 1 wins), -1 (player 2 wins), or 0 (a tie). Representing the game in strategic form (unlike extensive form) means that we do not assume that one of the players chooses his strategy before the other but that both players choose their strategies simultaneously.*

Example 3

	left	middle	right
up	1	4	5
down	3	2	6

Player 1 has two possible strategies, player 2 has 3 possible strategies.

Let us look at Example 3. What would player 1 do if he had to choose his strategy before player 2 chooses a strategy? If he were to choose "up" then player two would obviously choose "left" and player 1 would only gain 1. However, if player 1 chose "down" then player 2 would choose "middle" and player 1 would gain 3. Therefore, assuming that both players are rational it would be best for player 1 to choose "down". In general, player 1 would choose the strategy for which player 2 can cause the least damage to player 1. Formally, this means that player 1's strategy would be s_r^1 such that $r = argmax_i min_j a_{ij}$.

What would player 2 do in Example 3 if he had to choose his strategy before player 1? Obviously, player 2 would choose "left" as it guarantees that the most that player 1 could gain is 3 (while he can gain more if player 2 chooses one of the other two strategies). Generally speaking, player 2 would pick strategy s_s^2 such that $s = argmin_j max_i a_{ij}$.

The following simple lemma shows that the value player 1 achieves if he chooses a strategy first is at most the value player 2 can achieve if he chooses first.

Lemma 1 $max_i min_j a_{ij} \leq min_j max_i a_{ij}$.

Proof: It is enough to show that for every i_0, j_0 we have that $min_j a_{i_0 j} \leq max_i a_{i j_0}$. This is true because $min_j a_{i_0 j} \leq a_{i_0 j_0} \leq max_i a_{i j_0}$ ■

We are interested in determining when $max_i min_j a_{ij} = min_j max_i a_{ij}$. In other words, when does it not matter which of the two players chooses his strategy first? When dealing with *pure strategies* it is not always true that $max_i min_j a_{ij}$ equals $min_j max_i a_{ij}$. Von Neumann showed that an equality exists with *mixed strategies*. A mixed strategy is a distribution over pure strategies (what we so far referred to merely as strategies).

Definition 2 *A mixed strategy of a player i with m_i pure strategies is a vector $x = (x_1, ..., x_{m_i})$ in $\Re^{m_i}$ such that $\forall t\ x_t \geq 0$ and $\Sigma_t x_t = 1$.*

Let us now consider a zero-sum two-player game with mixed strategies. Let $x = (x_1, ..., x_n)$ be a mixed strategy of player 1 and $y = (y_1, ..., y_m)$ be a mixed strategy of player 2. What is the value of the game's outcome when 1 plays x and 2 plays y? This value is defined to be $\Sigma_i \Sigma_j x_i y_j a_{ij}$. An implicit assumption is that both players are *expected utility maximizers*. That is, we assume that the players only wish to maximize their expected utility and not the utility they get in some specific toss of random coins.

Let us now assume that player 2 chooses his strategy after player 1 chose x. As in the case of pure strategies, he would choose the mixed strategy that minimizes 1's gain. Notice that since every mixed strategy is a distribution on pure strategies 2 can simply choose the best pure strategy (if there is more than one best pure strategy he can choose any distribution over these strategies). Formally, 2 will choose the column j for which $\Sigma_i x_i a_{ij}$ is minimized. Similarly, if player 1 chooses his strategy after player 2 chose y then he would choose the row i such that $\Sigma_j y_j a_{ij}$ is maximized.

Therefore, if player 1 chooses a strategy before player 2 he would choose $x = argmax_s min_j \Sigma s_i a_{ij}$. Similarly, if player 2 had to choose before player 1 he would choose $y = argmin_r max_i r_j a_{ij}$. Von Neumann proved that when dealing with mixed strategies $max_s min_j \Sigma s_i a_{ij} = min_r max_i r_j a_{ij}$. That is, intuitively, with mixed strategies, it does not matter which player chooses a strategy first.

We want to design an algorithm that finds the mixed strategy that guarantees the highest gain for player 1 (assuming that player 2 is rational). The *input* for this algorithm would be the value matrix A, and the output would be a mixed strategy x such that $x = argmax_s min_j \Sigma s_i a_{ij}$. We will show that being able to solve this problem is precisely equivalent to being able to solve *linear programming problems*.

A *mathematical programming problem* deals with finding real values for parameters $x_1, ..., x_n$ that maximize some function $f(x_1, ..., x_n)$ such that a set of inequalities of the form $g_t(x_1, ..., x_n) \geq 0$ hold. We can formulate our problem as a mathematical programming problem by setting $f = min_j \Sigma_i x_i a_{ij}$ and requiring that $x_i \geq 0$ for all i and that $\Sigma_i x_i = 1$. If f and all g_t's are linear functions then the mathematical programming problem is called a linear programming problem (*LP problem*).

We shall convert our problem to an LP problem in the following way: We will maximize a parameter c such that $x_i \geq 0$ for all i and $\Sigma_i x_i = 1$ and require that for all j's $c \leq \Sigma_i x_i a_{ij}$. It is easy to show that this LP problem is equivalent to the original mathematical programming problem.