Algorithmic Game Theory March 2, 2016

Lecture 1

Lecturer: Michal Feldman *Scribe: Jacob Komarovski, Chen Amar, Itay Polack*

1 Administration

Instructor: Michal Feldman

TA: Israela Solomon, israela5@mail.tau.ac.il

Website: agttau-2016.wikidot.com

1.1 Course Requirements

Homework Assignment

- 30% of final grade.
- 3 assignments submitted in pairs
- Each question will be given a grade between 0 and 3.

Scribe Notes

- Scribe notes: 10% of final grade.
- Should be submitted in latex (the template can be found on the course site).

Project

- 60% of final grade. Will be submitted in groups.
- Can be either a research project, or a survey project.
- Research project is the preferred option. The students will be required to show a good understanding of the papers and articles, of the subjects covered in class and maybe even publish their research.

- A survey project: the students will be required to pick a few papers, read them and write a comprehensive sum-up. They need to present some added value, compare between different modules/methods.
- There will be a few milestones during the semester and the students will have the opportunity to meet with other staff members related to the subject of the project.

2 Algorithmic Game Theory - Introduction

The rise of algorithmic game theory is due to the rise of the internet. This is a mixture of economics and computer science.
Game theory tries to analyze situations from real life and involving people using strategies in order to maximize their own benefit.

Example:IBM is trying to implement an algorithm that will work correctly and will not take into effect human manipulations. Let's assume that an algorithm exists for regulating and controlling network traffic (TCP for example). When the network gets congested, TCP limits the rate. This protocol is implemented on all computers but each computer has its own priorities (someone may want to send a lot of data at a specific moment) which can lead to a situation where a computer won't follow the instructions of the protocol because it has its own incentives.

This area of studies tries to analyze and understand a given situation and turns to the field of economics where these questions have already been asked and researched for a long time. We will try and bring this expertise into the field of computer science.
Game theory is different from the decision theory. In decision theory, if we look at the question "should I take an umbrella?" we need to decide if we take the umbrella or not. There is an element (nature) that effects our benefit.
In game theory, there is extra layer - probability. Let's say that we need to get from Tel Aviv to Jerusalem. We are not the only ones trying to do it because there are a lot more people using this route so our decision is not the only thing effecting us. If everyone chose route A and we choose route B this will be good for us but not for long, because everyone else will figure it out and change course.
We will want to get to a "clean" model that tries to predict the eventual outcome. There are games without any knowledge, little knowledge or ones that have all the information we need.
There also exists a reverse impact. The field of economics is effected by computer science because of growing amounts of options for social interactions, the emergence of electronic trading websites such as EBay where there is a big selection of products and large competition. There are dependencies between situations (IPhone + IPhone earphones), interchangeable situations (IPhone 5 vs IPhone 6) etc.
In essence, this is a field of planning mechanisms –We will want to understand in advance

how to plan the game (mechanism) in such a way that even if the players do what's best for them we will still maximize our own benefit (utility). Add to this the computational factor and ask ourselves how we can best achieve our goal function. In this way we limit ourselves to the algorithms we want and allow us to achieve our goal.

This course has two parts:

1) Analysis of existing systems and price of anarchy. *Price of anarchy* is the ratio between an unorganized and organized system.

Example: Braess's Paradox (based on real life events)

Figure 1: Roads

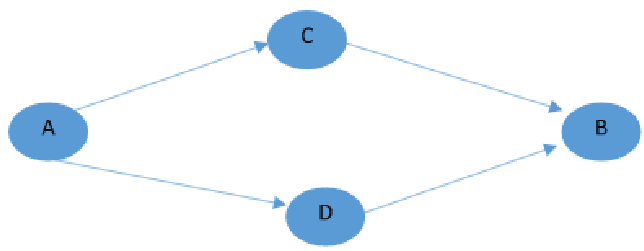

Each car needs to get from A to B. It needs to choose a route via C or D. Each road (edge) has a weight which is the time it takes to cross the road with a congestion of x , c(x).

$$C_{AC}(\text{x})=\text{x},\ C_{AD}(\text{x})=1,\ C_{CB}(\text{x})=1,\ C_{DB}(\text{x})=\text{x}$$

If the congestion is $\frac{1}{100}$ then it takes $\frac{hour}{100}$ time to cross the road.

The utility function is the average between the travels times and we will try to minimize it. In order to minimize the utility function we will send half of the cars via C and half via D. Definition: equilibrium —A stable state where there is no benefit in changing strategy. We will notice that the above solution is an equilibrium.

Now let's alter the graph:

Figure 2: Roads

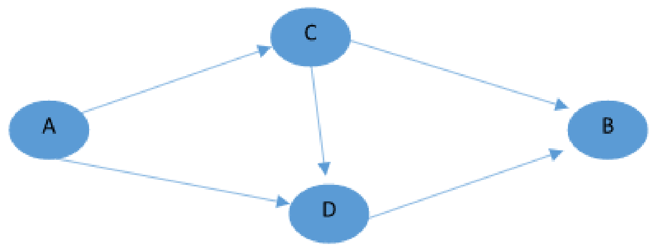

Where :

$$C_{AC}(\mathrm{x})=\mathrm{x},\ C_{AD}(\mathrm{x})=1,\ C_{CB}(\mathrm{x})=1,\ C_{DB}(\mathrm{x})=\mathrm{x},\ C_{CD}(\mathrm{x})=0$$

Now, every car going from A to C, C to D, D to B can shorten the time. The equilibrium is reached if all cars use the route ACDB which takes 2 hours (not shorter than the route ACB for example).

We have added another road and didn't shorten our travel time –This is Brauss's Paradox.

Another interesting view point for the paradox is through a physical model, demonstrated by a system of springs and weights (real experiments can be watched in Youtube).

2) Planning mechanism —How to produce a mechanism (game) that is the best for everyone. A lot of times there is a conflict between the planner of the game and the players.

3 Basics

3.1 Definitions

For start, we will take a few assumptions:

1. Games are **not cooperative**: each player tries to maximize their own benefit with no regard to other players' payoff.
2. Players have **full information** - they know all available strategies and payoffs of all players.
3. **Strategic game** - we ignore the time dimension: all players act simultaneously, and do not change their behavior between the turns.

A game of n players is defined by the following:

1. Set of strategies $S = \{S_1, S_2, \ldots, S_n\}$.
2. Utility functions. For each player i, let $u_i : S_1 \times S_2 \times \ldots \times S_n \to \mathbb{R}$ be their utility function.

Definition 1 *$S_{s_i \in S_i} = (s_1, \ldots, s_n)$ is a **strategy profile**, representing the strategies of all players. We will use the $S_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ notation to represent the strategies of all players except player i.*

Definition 2 *$s_i \in S_i$ is a **best response** to a strategy profile $s_{-i} \in S_{-i}$ if it achieves the best payoff to player i, given the other players' strategies are s_{-i}. Formally, strategy profile $s_i \in S_i$ is best response to $s_{-i} \in S_{-i}$ if $\forall_{s'_i \in S_i} u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$.*

Definition 3 *Strategy s_i **dominates** strategy s'_i if s_i achieves better or equal payoff for all strategies played by the other players. Formally, $\forall_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$.*

Definition 4 *s_i is a **dominating** strategy for player i if it dominates all other strategies for this player.*

Definition 5 *s_i is **dominated** if there exists a strategy s'_i for player i that dominates it.*

3.2 Nash Equilibrium in Pure Strategies

Definition 6 *$S = (s_1, \ldots, s_n)$ is called Nash Equilibrium if for each player i, s_i is the best response for s_{-i}.*

We will see examples for a few 2 players games. The games are represented by a matrix, where each cell $a_{i,j}$ is a vector with the utilities each player gets, given player 1 ("rows") employs strategy s_i, and player 2 ("columns") employs strategy s_j.

Rock-paper-scissors

In this popular game, rock beats scissors but lose to paper; Paper beats rock but lose to scissors; Scissors beats paper but loses to rock. When both players play the same strategy, it's a tie. We will describe this game using the table below.

Is there a dominating strategy for this game? The answer is no: each strategy only works against a single other strategy. There is no single strategy that guaranteed to give better results than all other strategies.

Table 1: Rock-Paper-Scissors matrix

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

The Prisoner Dilemma

Each player can choose to either cooperate (C) or defect (D). For each strategy set, the payoffs are as described in table below.

Table 2: Prisoner dilemma matrix

S	C	D
C	$(3,3)$	$(0,5)$
D	$(5,0)$	$(1,1)$

In this famous game, if both players cooperate, they both get a high award. If both defect, they are both penalized with low reward. If one player choose to cooperate while the other choose to defect, the defecting player is rewarded and the cooperating player is severely penalized.

In this case, there is a dominating strategy: D. Both sides are getting smaller payoff when taking this strategy, but choosing to cooperate (C) is an unstable status, because the other player will choose D in order to improve their payoff.

This dilemma fits other situations as well, such as - keeping clean environment. Everyone wants to enjoy clean environment (the results of everyone picking the cooperation strategy), but on the other hand - does not want to put the effort to keep clean (let someone else care). The end result - nobody is keeping clean, and everyone suffers dirty environment. Such situation is called "tragedy of the commons".

Battle of the Sexes

Each player can choose between going to the opera (O), or to a boxing match (B). Going together to either of the options benefits both players (with a small additional payoff to the player whose preference was chosen), while going to different events is a lose-lose situation.

Table 3: Battle of the Sexes matrix

S	B	O
B	$(3,2)$	$(0,0)$
O	$(0,0)$	$(2,3)$

For this game, a best-response strategy exists for each possible strategy, but there is no single dominating strategy.

Nash equilibrium exists, but has multiple sets of strategies: $\{< B, B >, < O, O >\}$.

3.3 Nash Equilibrium in Mixed Strategies

Observation 7 *Through the above examples we reach two important observations:*

1. *Nash equilibrium does not always exists (as in rock-paper-scissors).*
2. *Nash equilibrium is not necessarily unique (as in the battle of the sexes).*

We will define **mixed strategies**.

Definition 8 *Let* $\Delta(s) = \{x_1, x_2, \ldots, x_{|s|} | \forall j x_j \geq 0, \sum x_j = 1\}$ *be a set of probabilities over the set of strategies* S. $x^i = \{x_1^i, x_2^i, \ldots, x_{|s|}^i\} \in \Delta(S)$ *is a* ***mixed strategy*** *for player* i*, where* x_j^i *is the probability that player* i *picks strategy* s_j*.*

Following this new definition, we will have to re-define the utility function as well:

Definition 9 ***Utility function with mixed strategies*** *is an expected value over the the different probabilities:* $u_i(x^1, ..., x^n) = \mathbb{E}_{s_j \sim x_j}[u_i(s_1, \ldots, s_n)]$.

Notice a strong assumption we take here: the players only care about expected value, completely ignoring risk. This assumption is not valid in many real world situation, but we will assume it for now.

With the new definitions, we can now rethink the paper-rock-scissors game we discussed before. We can now achieve Nash equilibrium using the following strategy set (same for all players): $\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$. This is an equilibrium, as the expected value for each player is 0, and changing probabilities can only reduce this expected value. For example, if some player gives higher probability to the "rock" option, the other player can increase the probability of the "paper" strategy and achieve a higher expected value.

Definition 10 ***Nash equilibrium in mixed strategies*** *is a vector* $(x^1, \ldots, x^n)$ *such that for all* i*,* x^i *is the best response for* x^{-i}*.*

Notice that in an equilibrium, the expected value for each strategies must be the same, otherwise the player can increase the probability of the strategy that gives them a better payoff.

Definition 11 *In a* ***pure strategy****, one strategy has probability of* 1*, and the rest of strategies* 0*.*

Example 1 *We will calculate NE for the "battle of the sexes" game. Assume that the mixed strategy for player 1 is* $(p, 1-p)$*. Since the expected value must be equal for all strategies, the following equation must hold:* $3p = 2(1-p) = 2 - 2p \rightarrow 5p = 2 \rightarrow p = \frac{2}{5}$*. We found out that the mixed strategy for player 1 is* $(\frac{2}{5}, \frac{3}{5})$*. In the same way, the mixed strategy for player 2 is* $(\frac{3}{5}, \frac{2}{5})$*.*

Theorem 12 *For any finite game (meaning: finite number of players and finite number of strategies) in mixed strategies, Nash equilibrium exists.*

The proof is based on another theorem (that we will not prove here), named **Brouwer Fixed-Point Theorem**.

Theorem 13 ***Brouwer Fixed-Point Theorem:*** *Let* C *be a closed bounded convex set in* $\mathbb{R}^t$*, and let* $f : C \rightarrow C$ *be a continuous function, then there exists a point* $x \in C$ *such that* $f(x) = x$*.*

Example 2 *Let* $t = 1$*,* $C = [0, 1]$*, and* $f : [0, 1] \rightarrow [0, 1]$ *- a continuous function. We will define* $g(x) = f(x) - x$*. The following holds:* $g(0) = f(0) \geq 0$*,* $g(1) = f(1) - 1 \leq 0$*. From continuity of* g *and the intermediate value theorem, there must exists some* $x' \in C$ *such that* $g(x') = 0$*. For this* x'*,* $g(x') = f(x') - x' = 0 \rightarrow f(x') = x'$*.*

Proof for Nash Equilibrium Theorem

We will define a function $f : \{\Delta(s_1) \times \ldots \times \Delta(s_n)\} \rightarrow \{\Delta(s_1) \times \ldots \times \Delta(s_n)\}$. We will denote the best response for x^{-j} as $BR(x^{-j})$.

According to the fixed-point theorem, the function we defined $f : (x^1, \ldots, x^n) \rightarrow BR(x^{-1}), \ldots, BR(x^{-n})$ must have at least one point $X' = (x'^1, \ldots, x'^n)$ where $f(X') = X'$, and that is exactly an equilibrium.

However, f is neither continuous nor injective (there could be multiple best responses for a strategy) - we cannot use the fixed-point theorem and therefore this proof is incorrect.

We will denote x_i^j to be the probability of a pure strategy j in distribution x_i. We will define $C_i^j = u_i(j, x^{-i}) - u_i(x_i, x^{-i})$ as the difference between the utility of player i if they choose the pure strategy j instead of choosing the strategy x_i. Let $C_i^{j+} = \max(C_i^j, 0)$.

Now define: $\hat{x}_i^j = \frac{x_i^j + C_i^{j+}}{1 + \sum_j C_i^{j+}}$ (the denominator normalizes this expression to be a proper distribution). Let $f(x^1, \ldots, x^n) = (\hat{x}^i, \ldots, \hat{x}^n)$. We assume that the C_i^j is continuous, therefore f is continuous as well. According to the fixed-point theorem, there exist a fixed point x such that $\forall i x^i = \hat{x}^i$, which also implies $\forall j x_i^j = \hat{x}_i^j$.

To complete the proof, we must show that for all i, j, $C_i^j \leq 0$ on that fixed-point. Otherwise, it means player i can improve their payoff by shifting to a pure strategy. If the condition is met, the strategy cannot be improved and therefore equilibrium is achieved.

Assume by contradiction the existence of i, j such that $C_i^j > 0$. It means that $\hat{x}_i^j = \frac{x_i^j + C_i^{j+}}{1+\sum\limits_j C_i^{j+}} > 0$, because both $x_i^j + C_i^{j+} > 0$ and $1 + \sum\limits_j C_i^{j+} > 0$.

We assume x_i^j is a fixed point, so $\hat{x}_i^j = x_i^j$, implying that $x_i^j > 0$.

For all i, $\sum\limits_j C_i^j x_i^j = 0$ (proof left to the reader), therefore there must exist some j' such that $C_i^{j'} < 0, x_i^{j'} > 0$ to balance the sum. By definition, $C_i^{j+} = 0$. From the fixed-point assumption: $x_i^{j'} = \hat{x}_i^{j'} = \frac{x_i^j + 0}{1+\sum\limits_j C_i^{j+}} < x_i^{j'}$. We reached a contradiction. ■

We proved that NE in mixed strategies exists for any finite game, but can we actually compute it? In the general case, finding NE is considered a computationally hard problem - even with 2 players games! However, for certain cases, we can find it efficiently.

3.4 Zero Sum Games with 2 Players

Definition 14 *A zero sum game is a game where the sum of the payoffs for all players is always* 0.

For example, rock-paper-scissors is a 2-players zero-sum game. When describing a zero-sum game with 2 players, a utility vector is not required - we can simply write the utility for player 1, and the utility for player 2 is its negative.

An example game:

Table 4: Zero-sum 2 players example game

S	L	M	R
U	0	20	100
D	10	1	200

When player 1 ("rows") takes strategy i (U or D), and player 2 ("columns") takes strategy j (L, M, or R), the payoff for player 1 is a_{ij}, and $-a_{ij}$ for player 2.

We we now assume player 1 plays first, and player 2 second. Player 1 will always play D, otherwise player 2 can play L and achieve a non-negative payoff. Player 2 will choose M, in order to achieve the best possible payoff in this situation. Eventually, the payoff for player 1 would be 1, and -1 for player 2.

In general, player 1 wants to choose the strategy in which player 2 can inflict as little damage as possible, meaning $r = \arg\max_i\{\min_j a_{ij}\}$. Following the same rationale, if player 2 plays first, they'll choose $r = \arg\min_i\{\max\_ia_{ij}\}$.

Lemma 15 ***Minimax theorem:*** $\max_i \min_j a_{ij} \leq \min_j \max_i a_{ij}$

Immediate implication: with mixed strategies, the order of the turns does not matter - the players will take the same strategy sets.