Computational Aspects of Game Theory **Bertinoro Spring School 2011**

Lecture 1: Overview

Lecturer: Bruno Codenotti

We first give a brief informal introduction to Game Theory, Social Choice Theory, and Mechanism Deign, and then present some connections between these disciplines and Computer Science.

1.1 Introduction

Game Theory (GT from now) is a scientific discipline which deals with the behavior of *homo rationalis* in interactive situations. Homo rationalis is an imaginary species that always acts with a goal, and has the ability to compute whatever needed to achieve the given goal. Homo rationalis has a (distant) relative in the real world, *homo sapiens*, who instead is influenced by his subconscious, and not always has precise goals. GT aims at understanding *something* about homo sapiens by using homo rationalis as a (very crude) model.

Definition 1.1 (Common knowledge of rationality) *Not only we assume that each player is rational, but also that each player knows that everybody else is rational, and that each player knows that everybody else knows that everybody else is rational, etc. We say that the fact that all the players are rational is* common knowledge.

Example 1.2 (The Beauty Contest) *Consider a game, where each player has to write an integer number between* 0 *and* 100 *on a piece of paper. The winner(s) will be the player(s) who has written the number closer to the* $\frac{2}{3}$*rd of the mean of the numbers written down. We assume that rationality is common knowledge among the players.*

Any given player, who is not allowed to communicate with the other players, will argue as follows. The $\frac{2}{3}$*rd of the mean cannot exceed* 67*, and so nobody will write down a number greater than* 67*. Therefore the* $\frac{2}{3}$*rd of the mean cannot exceed* 44*, and nobody will write down a number greater than* 44*. By iterating this process, the rational player (who also assumes that everybody else is rational) will write down the number zero*[1].

GT aims at modeling and analyzing the nature of real-life phenomena produced by the interactions between *decision-makers*. Note that the notion of decision-maker (or agent, or player) does not necessarily model a single individual, but more in general an entity which makes decisions (a company, an association, a political party, etc). GT uses analytical/mathematical tools to analyze such phenomena.

Informally, a *game* consists of a set of interactions among a number of players which are in contention for a reward (or payoff). They have to make moves on which their payoff depends, and they have to follow certain *rules* while making moves. All the players behave *rationally*, e.g., they try to maximize their payoff (*rationality assumption*).

A game is a highly abstract description of strategic interactions in competitive situations, and does not specify the actual actions taken by the players, but rather it assumes their rational behavior. A *solution* to a game is a complete description of the *outcomes* that may arise.

There are two aspects of GT:

[1] Those interested in knowing why this game has been called The Beauty Contest can see Chapter 12 of [7].

- A *normative* aspect, which uses GT to give suggestions to the players. For instance, how to split a payoff so to satisfy certain fairness properties, how to negotiate an international treaty, etc.
- A *classification* aspect, where the focus is on subdividing the interactive situations in different types. For example, games where players can enforce agreements vs games where players cannot communicate, or games where the differences are based on the information available to the players, etc.

The most basic subdivision between games is determined by the availability or not of suitable mechanisms to enforce agreements among the players. This leads to two settings:

- *Non-cooperative games*, where the assumption is that the *primitive* operations are at the player's level. Here players do not cooperate, and act independently, without communicating. More precisely, the lack of mechanisms to enforce agreements makes communication useless.
- *Cooperative games*, characterized by primitive operations at the coalitional level. Here players can interact, communicate, form coalitions, and can make enforcing agreements.

1.2 Non-Cooperative Games

In a non-cooperative game, there is no external mechanism available to enforce agreements. Players act on the basis of existing incentives.

An important distinction is between *strategic* and *extensive form* games:

- A *strategic game* (or a *game in normal form*) is a game where all players (simultaneously) decide their plan of action once at the beginning of the play.
- An *extensive form game* specifies the possible order of events, and each player can decide her plan of action not only at the beginning of the play, but whenever she is called to make a *move.*

The fundamental notion behind the choices of the players is that of *strategy.* A strategy is a complete plan of action, that defines the set of moves a player will take in a given game. In this case we talk of pure (or deterministic) strategies.

We now consider two-player games in *normal form.* These games are described by a pair (A, B) of matrices, whose entries are the *payoffs* of the two players, called row and column players. $A = (a_{ij})$ is the payoff matrix of the row player, and $B = (b_{ij})$ is the payoff matrix of the column player.

The rows (resp. columns) of A and B are indexed by the row (resp. column) player's *pure strategies.*

The entry a_{ij} is the payoff to the row player, when she plays her i-th pure strategy and the opponent plays his j-th pure strategy. Similarly, b_{ij} is the payoff to the column player, when he plays his j-th pure strategy and the opponent plays her i-th pure strategy.

A *mixed strategy* is a probability distribution over the set of pure strategies which indicates how likely it is that each pure strategy is played. More precisely, in a mixed strategy a player associates to her i-th pure strategy a quantity p_i between 0 and 1, such that $\sum_i p_i = 1$, where the sum ranges over all pure strategies. In other words, both players pick the pure strategies at random according to some probability distribution of their choice.

Let us consider the game (A, B), where A and B are $m \times n$ matrices. In such a game the row player has m pure strategies, while the column player has n pure strategies. Let x (resp. y) be a mixed strategy of the row (resp. column) player. Strategy x is the m-tuple $x = (x_1, x_2, \ldots, x_m)$, where $x_i \geq 0$, and $\sum_{i=1}^m x_i = 1$. Similarly, $y = (y_1, y_2, \ldots, y_n)$, where $y_i \geq 0$, and $\sum_{i=1}^n y_i = 1$.

When the pair of mixed strategies x and y is played, the entry a_{ij} contributes to the expected payoff of the row player with weight $x_i y_j$. The expected payoff of the row player can be evaluated by adding up all the entries of A weighted by the corresponding entries of x and y, i.e., $\sum_{ij} x_i y_j a_{ij}$. This can be rewritten as $\sum_i x_i \sum_j a_{ij} y_j$, which can be expressed in matrix terms as[2] $x^T Ay$. Similarly, the expected payoff of the column player is $x^T By$.

Definition 1.3 (Stochastic vector) *A k-vector $z = (z_i)$ with real nonnegative entries such that $\sum_{i=1}^k z_i = 1$ is called a stochastic k-vector.*

Definition 1.4 (Nash Equilibrium) *A pair of mixed strategies (x, y) is in Nash equilibrium if $x^T Ay \geq x'^T Ay$, for all stochastic m-vectors x', and $x^T By \geq x^T By'$, for all stochastic n-vectors y'.*

We say that x (resp. y) is a Nash equilibrium strategy *for the row (resp. column) player.*

The set of indices such that $x_i > 0$ (resp. $y_i > 0$) is called the support *of the Nash equilibrium strategy x (resp. y).*

In a non-cooperative game where players use their pure strategies, a Nash equilibrium is not guaranteed to exist. However, if one assumes that players can resort to mixed strategies, then a Nash equilibrium always exists (by the celebrated Nash Theorem [12, 13]).

1.3 Cooperative games

Cooperative game theory asks questions about the agreements that players should reach. Unlike the theory of non-cooperative games, cooperative game theory does not specify a game by strategies, moves, actions, etc., but rather focuses on the coalitional form, and, in particular, on the payoff opportunities available to each coalition.

A cooperative game (or coalitional game) in characteristic form is defined by a set $N = \{1, 2, \ldots, n\}$ of players and by a function $v(\cdot)$ which maps the set of all possible coalitions into the real numbers. For any coalition $S \subseteq N$, $v(S)$ is the (monetary) worth of S. We can think of $v(S)$ as the payoff that the players in S can guarantee for themselves, independently of the other players. So $v(\{i\})$ denotes the payoff that player i can obtain by playing in a non cooperative way. In this sense one can view a cooperative game as starting *after the end* of an associated non cooperative game which gave a payoff of $v(\{i\})$ to player i, $i = 1, 2, \ldots, n$. The value $v(N)$ represents the total payoff available to the *grand coalition* of all the players.

The most natural counterpart to the Nash equilibrium, in the context of coalitional games, is the notion of *core*. An *imputation* to the players is a vector $\mathbf{x} = (x_1, x_2, \ldots, x_n)$, where x_i is the payoff to player i, and such that $x_i \geq v(\{i\}), \forall i \in N$, and $\sum_{i \in N} x_i = v(N)$. The core describes the set of all those imputations from which no coalition of players has an incentive to unilaterally deviate. Indeed an imputation x is in the core if $x(S) = \sum_{i \in S} x_i \geq v(S)$, for all $S \subseteq N$, with equality for $S = N$. In other words, a core imputation is feasible, i.e., $x(N) = v(N)$, and it assigns to each possible coalition at least as much as the coalition could guarantee for itself. This guarantees that there is no coalition S such that $\sum_{i \in S} x_i < v(S)$. If this were the case, players in S would have an incentive not to accept the imputation and rather divide among themselves the amount $v(S)$. Given a cooperative game (N, v), we will denote its core by $Core(N, v)$.

[2] We use the notation x^T to denote the transpose of vector x.

The trouble with the notion of the core is that it might be empty, as we see from Example 1.5 below.

Example 1.5 *Three persons have to move a weight from one location to another. The total reward for this job is one dollar. Any two of the three players can do the job, but none of them can do it alone. They can collaborate, and make all kinds of internal arrangements about the subdivision of the payoff. Each of them tries to maximize her own payoff.*

We can model this situation by introducing a function v defined for each subset S of the players:

$$\mathrm{v}(S) = \begin{cases} 0 & \text{if } |S| \leq 1 \\ 1 & \text{if } |S| \geq 2 \end{cases}$$

$v(S)$ can be thought of as the payoff which no one can prevent the players in coalition S from obtaining. The game can be described by the pair $< N, v(\cdot) >$, where $N = \{1, 2, 3\}$ is the set of players.

The goal is to find a core imputation, i.e., an imputation $x_1 \geq 0$, $x_2 \geq 0$, and $x_3 \geq 0$ such that

$$\begin{cases} x_1 + x_2 \geq 1 \\ x_1 + x_3 \geq 1 \\ x_2 + x_3 \geq 1 \\ x_1 + x_2 + x_3 = 1 \end{cases}$$

If we add up the first three inequalities, we get $x_1+x_2+x_3 \geq \frac{3}{2}$, which is clearly incompatible with $x_1+x_2+x_3 = 1$. Therefore the core of this game is empty, i.e., there is no way to prevent two of the players to deviate from any imputation and divide the reward between themselves.

An intuitive explanation for the emptiness of the core is the following. If we look, e.g., at the allocation $(\frac{1}{2}, \frac{1}{2}, 0)$, we can see that the third player can destabilize it by suggesting to the second player the alternative allocation $(0, \frac{3}{4}, \frac{1}{4})$, for which both player 2 and 3 are better off.

The possible emptiness of the core is one of the reasons why the core is far from being the only solution concept adopted in coalitional games. Other notions include *stable sets*, the *kernel*, the *bargaining set*, the *Shapley value*. They all suggest, with different subtleties, some reasonable way of distributing the payoff available to the grand coalition.

Conceptually, the core can be viewed as a *first approximation*, with other solution concepts overcoming its shortcomings in different ways.

1.4 Markets

In a *strategic game*, each player takes decisions which depend on the strategies available to the other players. The market setting differs from the scenario of games, because the *equilibrium prices* – see below – have the "decentralizing" effect of making the strategic decisions of the economic agents independent. However there is a natural interplay between *market equilibria* and *Nash equilibria* for games: one of the very first proofs of existence of an economic equilibrium is built upon the existence of a Nash equilibrium in an associated abstract game. The actors in this game are the economic agents and an extra player, the market, whose strategy set coincides with the prices.

In 1874, Walras published the famous "Elements of Pure Economics", where he describes a model for the state of an economic system in terms of demand and supply, and expresses the *supply equal demand* equilibrium conditions [23]. In 1954, Nobel laureates Arrow and Debreu proved the existence of an equilibrium under mild assumptions [2].

We describe a model of the so-called *exchange economy*, an important special case of the model considered by Arrow and Debreu [2]. The more general one, which we will call the *Arrow-Debreu model*, includes the production of goods.

Let us consider m economic agents which represent traders of n goods. Let $\mathbf{R}_+^n$ denote the subset of $\mathbf{R}^n$ with all nonnegative coordinates. The j-th coordinate in $\mathbf{R}^n$ will represent good j. Each trader i has a *utility function* $u_i : \mathbf{R}_+^n \to \mathbf{R}_+$, which represents her preferences for the different bundles of goods, and an initial endowment of goods $w_i = (w_{i1}, \ldots, w_{in}) \in \mathbf{R}_+^n$. At given prices $\pi \in \mathbf{R}_+^n$, trader i will sell her endowment, and get the bundle of goods $x_i = (x_{i1}, \ldots, x_{in}) \in \mathbf{R}_+^n$ which maximizes $u_i(x)$ subject to the budget constraint[3] $\pi \cdot x \leq \pi \cdot w_i$.

An equilibrium is a vector of prices $\pi = (\pi_1, \ldots, \pi_n) \in \mathbf{R}_+^n$ at which, for each trader i, there is a bundle $\bar{x}_i = (\bar{x}_{i1}, \ldots, \bar{x}_{in}) \in \mathbf{R}_+^n$ of goods such that the following two conditions hold:

1. For each good j, $\sum_i \bar{x}_{ij} \leq \sum_i w_{ij}$.
2. For each trader i, the vector $\bar{x}_i$ maximizes $u_i(x)$ subject to the constraints $\pi \cdot x \leq \pi \cdot w_i$ and $x \in \mathbf{R}_+^n$.

1.5 Social Choice Theory

The theory of social choice deals with the decisions of a society obtained by aggregating the *preferences* of the individuals. The typical questions addressed involve (i) evaluating the *social decisions* in the presence of diverging interests of the individuals, and (ii) giving proper recognition to the preferences of the individuals. The milestones of social choice theory are impossibility results. The first of these is known as *Condorcet Paradox* (1785).

Consider an election with three candidates (A, B, and C), and three voters (1,2, and 3). Assume that the three voters rank the candidates in the following order

1. A B C
2. B C A
3. C A B

Who should be the winner of this election? If we perform a sequence of *pairwise majority* comparisons between the candidates, we notice that candidate A wins by majority against B, and so does B against C. Since in turn C wins against A, we can see that the preferences of the voters give rise to a cyclic behavior, which makes it problematic to choose a winner.

This paradox shows that a method of aggregating individual preferences in a social choice (in this case determining the winner of an election) by means of pairwise majority voting cannot work.

Arrow has shown that the difficulty pointed out by Condorcet Paradox has a general flavor, by proving a quite general *Impossibility Theorem* [1]. This Theorem states that if the aggregation rule has to (i) be defined over all possible preferences of individuals (*unrestricted domain*), (ii) provide decisions which reflect the preferences of individuals, if these are unanimous (*Pareto-efficiency*), and (iii) satisfy the property that the

[3] Given two vectors x and y, $x \cdot y$ denotes their inner product.

social preference over any two alternatives depends only on individual preferences over these two alternatives (*independence of irrelevant alternatives*), then the aggregation either lacks some rationality conditions or it is a *dictatorship*, i.e., the social preferences coincide with the preferences of a single individual.

The main message that we receive from these negative results is that social choice theory has to deal with unavoidable underlying difficulties.

In order to cope with these difficulties, one has to either restrict the domain of the individual preferences or radically change the model. The quest for "escape routes" takes us into the domain of Mechanism Design, a discipline whose objective is to implement desired social choices in strategic settings.

1.6 Mechanism Design

The theory of Mechanism Design can be viewed as the *engineering* side of Economic Theory. Much of the work in Economic Theory has to do with existing institutions, and the theoretician wants to explain or make predictions about the outcomes of these institutions. For example, in the context of GT, we fix a game, and analyze the set of possible outcomes. In the theory of Mechanism Design, the direction of investigation is inverted: one starts by identifying a desired outcome, or social goal, and then asks if there exists an institution (the *mechanism*) that makes it possible to reach the goal. For example, we can consider a set of outcomes, and try to come up with a game for which this set is the set of its equilibria. The challenge here is to design a game so that the actions of the individual players, despite their selfishness, lead to prescribed outcomes.

Mechanism design attempts to escape from the impossibility results of Social Choice Theory using several different modifications of the model. One of these is the addition of *money*. Money works as a yardstick that allows us to measure how much a given alternative is preferred to another. In a world with money, mechanisms will not only choose a social alternative, but will also determine monetary payments to be made by the different players.

1.7 Game Theory and Computer Science

One of the main motivations behind the growing interest in Game Theory within the computer science community is the Internet. The Internet is a computational artifact which emerged from the combined efforts of many different designers. The outcome of these efforts is an extremely large and complex object. As a result, it is reasonable to analyze features and properties of the Internet as one would with complex natural objects. Game Theory seems to be the right tool for this kind of investigation, since the entities involved in the Internet are locally optimizing agents interacting with one another with different levels of collaboration and competition.

Two different scenarios characterize the interplay between GT (or, more in general, Economic Theory) and Computer Science:

1. Economic Theory for Computer Science.
2. Computer Science for Economic Theory.

Economic Theory for Computer Science. Here we want to design algorithms (or protocols) in environments with multiple owners, where the algorithms have to cope with their potentially strategic (selfish) behavior.

The focus of a line of research has been on algorithmic problems in distributed settings where the agents cannot be trusted to follow a prescribed protocol, but rather their own self-interest [6, 14]. As such agents can manipulate the protocol, the designer should ensure that the agents' interests are best served by a correct behavior. The algorithm is typically endowed with a system of "payments" to the agents, chosen so as to motivate them to act as the designer wishes. This research area has been called *computational mechanism design*.

Computer Science for Economic Theory. Here we want to understand the computational complexity of the key game-theoretic parameters, design algorithms for their computation, and, on the most practical front, develop software for economic interactions, in particular for those occurring on the Internet.

A great deal of research has been done on the development of general methods for the computation of equilibria. Classical examples are the algorithms of Scarf and some coauthors for the computation of market equilibria [19], and the algorithm of Lemke and Howson for the computation of Nash equilibria [8, 20]. While these algorithms are very general, their running times are not polynomially bounded.

Over the last few years, the problem of computing equilibria has received a significant attention within the theoretical computer science community. In a short span of time, exciting developments have led to hardness results, the most remarkable being the proof of the hardness of two player games [5], as well as polynomial time algorithms for computing equilibria in several special cases.

Bibliographic notes

Aumann wrote an overview article [3] which gives a precise idea of the main goals of GT, and points out some of the crucial questions faced by game theorists.

The book by Osborne and Rubinstein [16] gives a well written presentation of both non-cooperative and cooperative game theory. The book by Mas-Colell et al. [10] is a very good reference for the market setting, as well as for social choice theory and mechanism design. Arrow's Impossibility theorem has been proved in [1]. For the interplay between Game Theory and Computer Science, one can start with the 1994 survey by Linial [9], and then see the handbook which was published in 2007 [15].

References

[1] K. Arrow, A Difficulty in the Concept of Social Welfare, Journal of Political Economy 58(4) (August, 1950), pp. 328346.

[2] K. Arrow, G. Debreu. Existence of an equilibrium for a competitive economy. Econometrica, 22:265–290, 1954.

[3] R. J. Aumann, What is game theory trying to accomplish? in K.Arrow and S. Honkapohja (eds.), Frontiers of Economics, Basil Blackwell, Oxford (1985).

[4] D. Black. Theory of Committees and Elections. Cambridge University Press, Cambridge, 1958.

[5] X. Chen, X. Deng, Settling the Complexity of Two-Player Nash Equilibrium, Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, pp. 261 - 272 (2006).

[6] R. K. Dash, N. R. Jennings, D. C. Parkes, Computational- Mechanism Design: A Call to Arms. IEEE Intelligent Systems, Volume 18, Issue 6 (November 2003), pp. 40-47 (2004).

[7] J. Keynes, The General theory of Employment, Interest and Money. Macmillan,London (1936).

[8] C.E. Lemke, J.T. Howson, Equilibrium points in bimatrix games, Journal of the Society for Industrial and Applied Mathematics 12, pp. 413- 423 (1964).

[9] N. Linial, Game-Theoretic Aspects of Computer Science, in Handbook of Game Theory with Economic Applications, vol. II, chapter 38, (R. J. Aumann and S. Hart eds.) North Holland, 1994, pp. 1340-1395.

[10] A. Mas-Colell, M.D. Whinston, J.R. Green, Microeconomic Theory, Oxford University Press (1995).

[11] R.B. Myerson, Game Theory: Analysis of Conflict, Harvard University Press, Cambridge, 1991.

[12] J. Nash, Equilibrium points in N-Person Games, 1950, Proceedings of NAS.

[13] J. Nash, Non-Cooperative Games, 1951, Annals of Mathematics.

[14] N. Nisan, Algorithms for Selfish Agents. In Proc. 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS'99), pp.1–15, 1999.

[15] N. Nisan, T. Roughgarden, E. Tardos, V. V. Vazirani, Editors, Algorithmic Game Theory, Cambridge University Press, September 2007.

[16] M. J. Osborne, A. Rubinstein. A Course in Game Theory. The MIT Press: Cambridge, MA, 1994.

[17] C. H. Papadimitriou, Algorithms, Games, and the Internet, STOC 2001 (extended abstract in ICALP 2001).

[18] C. H. Papadimitriou, Learning the Internet. Colt 2002. LNAI 2375, p.396 (2002).

[19] H. Scarf, T. Hansen. The Computation of Economic Equilibria. Cowles Foundation Monograph No. 24., New Haven: Yale University Press, 1973.

[20] L.S. Shapley, A note on the Lemke-Howson Algorithm, Mathematical Programming Study 1: Pivoting and Extensions, 175-189 (1974).

[21] L.S. Shapley, M. Shubik. A Method of Evaluating the Distribution of Power in a Committee System. American Political Science Review 48(3): pp. 787-792 (1954).

[22] J. von Neumann, O. Morgenstern. Theory of Games and Economic Behavior. Princeton Univ. Press (1944).

[23] L. Walras, Éléments d'économie politique pure, ou théorie de la richesse sociale (Elements of Pure Economics, or the theory of social wealth), 1874. (1899, 4th ed.; 1926, rev ed., 1954, Engl. transl. Homewood, Ill.: Richard Irwin.)