

Algorithmic Game Theory and Applications

Richard Mayr

U. of Edinburgh, UK

Fall Term 2014

Basic Course Information

- **Lecturer:** Richard Mayr; Office: IF 4.11; Phone: 650 5130.
- **Lecture times:** Monday & Thursday, 16:10-17:00; Lecture Theatre 4, 7 Bristo Square.
Tutorials: TBA. (starts in week 3).
- **Course page** (with lecture notes/reading list):
<http://www.inf.ed.ac.uk/teaching/courses/agta/>
- No required textbook. Course based on lecture notes + assigned readings. Some reference texts:
 - ▶ M. Maschler, E. Solan, & S. Zamir, Game Theory, 2013.
 - ▶ M. Osborne and A. Rubinstein, A Course in Game Theory, 1994.
 - ▶ R. Myerson, Game Theory: Analysis of conflict, 1991.
 - ▶ A. Mas-Colell, M. Whinston, and J. Green, Microeconomic Theory, 1995.
 - ▶ N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani (editors), Algorithmic Game Theory, 2007. (Book available for free online.)
 - ▶ Y. Shoham & K. Leyton-Brown, Essentials of Game theory, 2008. and Multiagent systems: algorithmic, game-theoretic, and logical foundations, 2009.
 - ▶ V. Chvátal, Linear Programming, 1980.

What is Game Theory?

A general and vague definition:

“Game Theory is the formal study of interaction between ‘goal-oriented’ ‘agents’ (or ‘players’), and the strategic scenarios that arise in such settings.”

What is *Algorithmic Game Theory*?

“Concerned with the computational questions that arise in game theory, and that enlighten game theory. In particular, questions about finding efficient algorithms to ‘solve’ games.”

These vague sentences are best illustrated by looking at examples.

A simple 2-person game: Rock-Paper-Scissors

		Player II		
		Rock	Paper	Scissors
Player I	Rock	0	-1	1
	Paper	1	0	-1
	Scissors	-1	1	0

- This is a “zero-sum” game: whatever Player I wins, Player II loses, and vice versa.
- What is an “optimal strategy” in this game?
- How do we compute such “optimal strategies” for 2-person zero-sum games?

A non-zero-sum 2-person game: Prisoner's Dilemma

Player II

		Cooperate	Defect
Player I	Cooperate	2, 2	0, 3
	Defect	3, 0	1, 1

- For both players Defection is a “*Dominant Strategy*” (i.e., regardless of what the other player does, you’re better off Defecting).
- But if they both Cooperate, they would both be better off.
- Game theorists/Economists worry about this kind of situation as a real problem for society.
- Often, there are no “dominant strategies”. What does it mean to “solve” such games?

Nash Equilibria

- A Nash Equilibrium is a pair (n-tuple) of strategies for the 2 players (n players) such that no player can benefit by unilaterally deviating from its strategy.
- **Nash's Theorem:** Every (finite) game has a mixed (i.e., randomized) Nash equilibrium.
- **Example 1:** The pair of dominant strategies (Defect, Defect) is a pure Nash Equilibrium in the Prisoner's Dilemma game. (In fact, it is the only Nash Equilibrium.)
In general, there may be many Nash equilibria, none of which may be pure.
- **Example 2:** In Rock-Paper-Scissors, the pair of *mixed* strategies: ((R=1/3, P=1/3,S=1/3), (R=1/3, P=1/3,S=1/3)) is a Nash Equilibrium. (And, we will learn, it is also a minimax solution to this zero-sum game. The "minimax value" is 0, as it must be because the game is "symmetric" .)
- **Question:** How do we compute a Nash Equilibrium for a given game?

Multiple equilibria

- Many games have > 1 Nash equilibrium.
Example: A “Coordination Game”:

		Player II	
		A	B
Player I	A	<div style="display: flex; justify-content: space-between;"> 2 2 </div>	<div style="display: flex; justify-content: space-between;"> 0 0 </div>
	B	<div style="display: flex; justify-content: space-between;"> 0 0 </div>	<div style="display: flex; justify-content: space-between;"> 1 1 </div>

- There are two **pure** Nash Equilibria:
 (A, A) and (B, B) .

Are there any other NEs?


Yes, there's one other *mixed* (randomized) NE.

Games in “Extensive Form”

So far, we have only seen games in “strategic form” (also called “normal form”), where all players choose their strategy simultaneously (independently).

What if, as is often the case, the game is played by a sequence of moves over time? (Think, e.g., Chess.)

Consider the following 2-person game tree:




- How do we analyze and compute “solutions” to extensive form games?
- What is the relationship to strategic form games?

chance, and information

Some tree nodes may be chance (probabilistic) nodes, controlled by neither player. (Poker, Backgammon.)

Also, a player may not be able to distinguish between several of its “positions” or “nodes”, because not all *information* is available to it. (Think Poker, with opponent’s cards hidden.) Whatever move a player employs at a node must be employed at all nodes in the same “information set”.

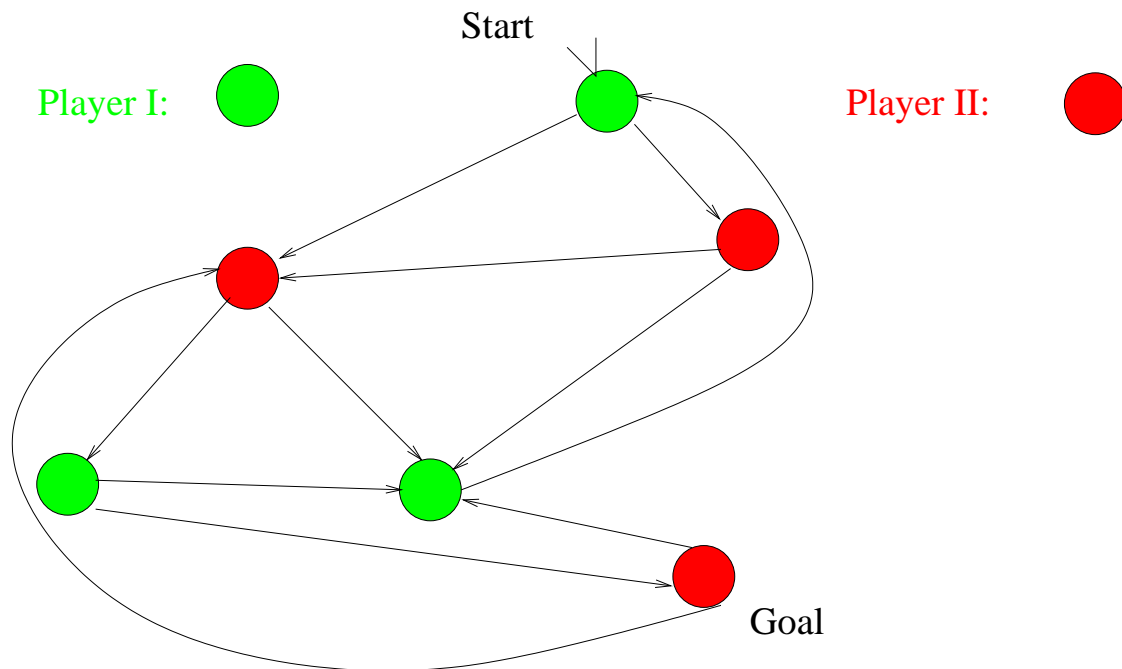


A game where every information set has only 1 node is called a game of perfect information.

Theorem A finite n-person extensive game of perfect information has an “equilibrium in pure strategies”.

Again, how do we compute solutions to such games?

Extensive games on Graphs



- Does Player I have a strategy to “force” the play to reach the “Goal”?
- Such games have lots of applications in, e.g., the analysis of programs in an “adversarial” input environment.
- Again, how do we compute winning strategies in such games?
- What if some nodes are chance nodes?

Mechanism Design

Suppose you are the game designer. How would you design the game so that the “solutions” will satisfy certain “global objectives”?

- **Example:** Auctions: (think EBay, or Google Ads)
Think of an auction as a multiplayer game between several bidders. If you are the auctioneer, how could you design the auction rules so that, for every bidder, bidding the maximum that an item is worth to them will be a “dominant strategy”?

One answer: Vickrey auctions: second price, sealed bid auctions.

- How would you design protocols (such as network protocols), to encourage “cooperation” (e.g., diminish congestion)?
- Many computational questions arise in the study of “good” mechanisms for various goals.
- This is an extremely hot topic of research (we will only get to scratch its surface).

But why study this stuff?

GT is a core foundation of mathematical economics. But what does it have to do with Computer Science? More than you might think:

GT ideas have played an important role in CS:

- Games in AI: modeling “rational agents” and their interactions. (Similar to Econ. view.)
- Games in Modeling and analysis of reactive systems: computer-aided verification: formulations of model checking via games, program inputs viewed “adversarially”, etc.
- Games in Distributed and Fault-tolerant computing: e.g., Byzantine agreement.
- Games in Algorithms: several GT problems have a very interesting algorithmic status (e.g., in NP, but not known to be NP-complete, etc).
- Games in Computational Complexity: Many computational complexity classes are definable in terms of games: Alternation, Arthur-Merlin

games, the Polynomial Hierarchy, etc..

Boolean circuits, a core model of computation, can be viewed as games (between AND and OR).

- Games in Logic in CS: GT characterizations of logics, including modal and temporal logics, and logics that capture computational complexity classes (Ehrenfeucht-Fraïssé games).

More recently:

- Games, the Internet, and E-commerce: An extremely active research area at the intersection of CS and Economics.

Basic idea: “The internet is a HUGE experiment in interaction between agents (both human and automated)”.

How do we set up the rules of this game to harness “socially optimal” results?

I hope you are convinced: knowledge of the principles and algorithms of game theory will be useful to you for carrying on future work in many CS disciplines.

Ok, let's get started

Definition A strategic form game Γ , with n players, consists of:

1. A set $N = \{1, \dots, n\}$ of players.
2. For each $i \in N$, a set S_i of (*pure*) strategies.
Let $S = S_1 \times S_2 \times \dots \times S_n$ be the set of possible combinations of (pure) strategies.
3. For each $i \in N$, a *payoff (utility) function* $u_i : S \mapsto \mathbb{R}$, describes the payoff $u_i(s_1, \dots, s_n)$ to player i under each combination of strategies.

(Each player prefers to maximize its own payoff.)

Definition A zero-sum game Γ , is one in which for all $s = (s_1, \dots, s_n) \in S$,

$$u_1(s) + u_2(s) + \dots + u_n(s) = 0$$

Some “food for thought”

At the end of every lecture, I’ll either give you HW, or I’ll try to give you a “food for thought” problem, not intended for you to write and hand in, but simply for you to think about and answer on your own.

Food for Thought:

(the “guess half the average game”)

Consider a strategic-form game Γ with n -players. Each player has to guess a whole number from 1 to 1000. The player who guesses a number that is closest to half of the average guess of all players wins a payoff of 1. All other players get a payoff of 0. (If there are ties for who is closest, all who are closest get payoff 1.)

Question: What would your strategy be in such a game?

Question: What is a “Nash Equilibrium” of such a game?