A4M33MAS - Multiagent Systems
Introduction to Game Theory

Michal Pechoucek & Michal Jakob

Department of Computer Science
Czech Technical University in Prague

OTEVŘENÁ INFORMATIKA

In parts based on Kevin Leyton-Brown: Foundations of Multiagent Systems an introduction to algorithmic game theory, mechanism design and auctions

Game Theory

- Game theory is the study of strategic decision making, the study of mathematical models of conflict and cooperation between intelligent rational decision-makers, interactive decision theory
- Given the *rule of the game*, game theory studies strategic behaviour of the agents in the form of a mixed/pure strategy (e.g. optimality, stability)
- Given the *strategic behavior of the agents*, mechanism design (reverse game theory) studies(designs) the rule of games with respect to a specific outcome of the game

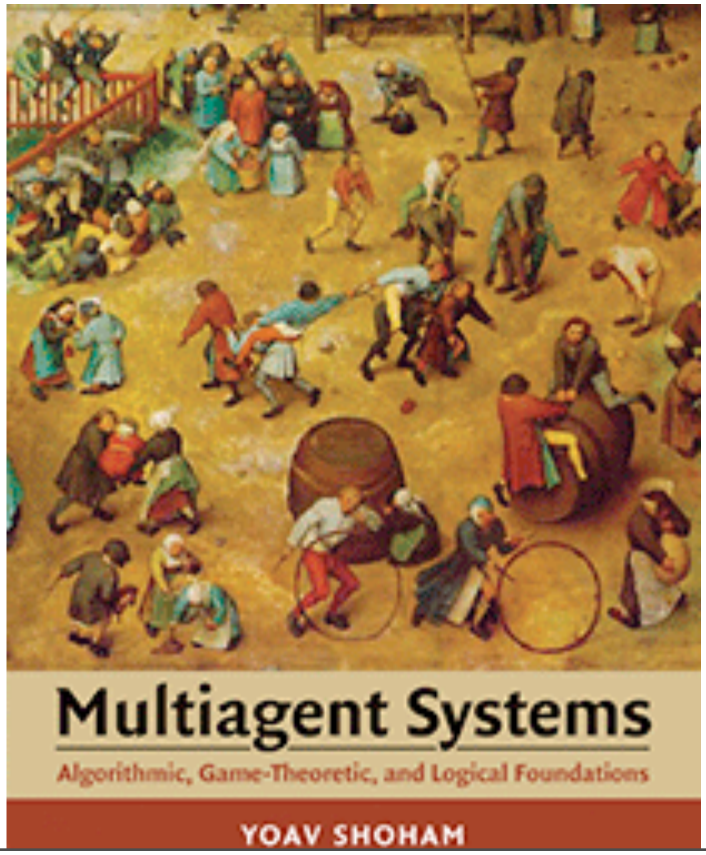

Yoav Shoham, Kevin Leyton-Brown, Multiagent Systems:
Algorithmic, Game-Theoretic, and Logical Foundations
Cambridge University Press, 2009
http://www.masfoundations.org

Types of Games

- Cooperative or non-cooperative

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games)

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games) Combinatorial games

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games) Combinatorial games
- Infinitely long games

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games) Combinatorial games
- Infinitely long games
- Discrete and continuous games, differential games

TCP Backoff Game

TCP Backoff Game

- Consider this situation as a two-player game:
  - both use a correct implementation: both get 1 ms delay
  - one correct, one defective: 4 ms delay for correct, 0 ms for defective
  - both defective: both get a 3 ms delay.

TCP Backoff Game

- Consider this situation as a two-player game:
  - both use a correct implementation: both get 1 ms delay
  - one correct, one defective: 4 ms delay for correct, 0 ms for defective
  - both defective: both get a 3 ms delay.

- Questions:
  - What action should a player of the game take?
  - Would all users behave the same in this scenario?
  - What global patterns of behaviour should the system designer expect?
  - Under what changes to the delay numbers would behavior be the same?
  - What effect would communication have?
  - Repetitions? (finite? infinite?)
  - Does it matter if I believe that my opponent is rational?

Game definition

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n players, indexed by i
  - $A = A_1 \times \ldots \times A_n$, where A_i is the action set for player i
    - $a \in A$ is an action profile, and so A is the space of action profiles
  - $u = \langle u_1, \ldots, u_n \rangle$, a utility function for each player, where $u_i : A \mapsto \mathbb{R}$

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Game definition

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Game definition

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Game definition

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

	C	*D*
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Other Games: Coordination Games

	Left	Right
Left	1	0
Right	0	1

driving side

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

battle of sexes

Other Games: Coordination Games

Players have exactly the same interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j, u_i(a) = u_j(a)$
- we often write such games with a single payoff per cell
- why are such games "noncooperative"?

Other Games: Prisoners Dilemma

	B_C	B_D
A_C	1 ,1	5, 0
A_D	0, 5	3, 3

$$(A_d, B_c)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_c, B_d)^5$$

$$(A_c, B_d)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_d, B_c)^5$$

Other Games: Prisoners Dilemma

	B_C	B_D
A_C	a, a	b, c
A_D	c, b	d, d

any game where $c \succeq a \succeq d \succeq b$

Other Games: Matching Pennies

	Heads	Tails
Heads	1 , -1	-1 , 1
Tails	-1 , 1	1 , -1

Other Games: Matching Pennies

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Other Games: Matching Pennies

	Heads	Tails
Heads	1 , -1	-1 , 1
Tails	-1 , 1	1 , -1

	Heads	Tails
Heads	1	-1
Tails	-1	1

Other Games: Rock-paper-scissors

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

Properties of the games

01

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n players, indexed by i
  - $A = A_1 \times \ldots \times A_n$, where A_i is the action set for player i
    - $a \in A$ is an action profile, and so A is the space of action profiles
  - $u = \langle u_1, \ldots, u_n \rangle$, a utility function for each player, where $u_i : A \mapsto \mathbb{R}$

- strategy s_i refers to a decision (about action choice) at each stage of the game that the agent i makes and which leads to an outcome
- outcome is the set of possible states resulting from agent's decision making
- strategy profile refers to the set of strategies played by the agents. Set of strategy profiles: $S = S_1 \times \ldots \times S_n$.

Properties of the games

- Social welfare (Collective utility):

$$U(a) = \sum_{\forall i} u_i(a_i)$$

- Cooperative agents choose such a_i that maximizes $U(a)$
- Self-interested (*individually rational*) agents choose such a_i that maximizes $u_i(a_i)$

- When designing a multiagent system designers worry about:
  - individual rationalty of each agent
  - social welfare and welfare efficiency
  - stability of the strategy (action) profile

Pareto Efficiency

- Pareto Efficiency:
  - action (strategy) profile is Pareto optimal if there is no other action that at least one agent is better off and no other agent is worse off than in the given profile

- Dominance:
  - measure comparing two strategies. b dominates weakly a as follows:

$$a \preceq b \text{ iff } \forall i : \; u_i(a_i) \leq u_i(b_i)$$

  - dominant strategy: strategy that is not dominated by any other strategy

Pareto efficient strategy is such a strategy that is not weakly dominated by any other strategy

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

Pareto Efficiency

	C	D
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Nash Equilibrium

01

- If you know what overyone else was going to do, it would easy to pick your own actions
- Let $a_i = \langle a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n \rangle$. now $a = (a_{-i}, a_i)$

Definition (Best Response)

$a_i^* \in BR(a_{-i})$ iff $\forall a_i \in A_i, u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$

Definition (Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff

$\forall i, a_i \in BR(a_{-i})$

Idea: look for stable action profiles.

Nash Equilibrium

Nash Equilibrium

- Nash equilibrium, is a set of strategies, one for each player, such that no player has incentive to unilaterally change her action. Players are in equilibrium if a change in strategies by any one of them would lead that player to earn less than if she remained with her current strategy.

Nash Equilibrium

- Nash equilibrium, is a set of strategies, one for each player, such that no player has incentive to unilaterally change her action. Players are in equilibrium if a change in strategies by any one of them would lead that player to earn less than if she remained with her current strategy.
- Strong Nash Equilibrium is such an equilibrium that is stable against deviations by cooperation.

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Strong Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Strong Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Prisoners Dilemma: PE, NE

	B_C	B_D
A_C	$1, 1$	$5, 0$
A_D	$0, 5$	$3, 3$

$$\xi_A = (A_d, B_c)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_c, B_d)^5$$

$$\xi_B = (A_c, B_d)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_d, B_c)^5$$

Prisoners Dilemma: PE, NE

PE

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

$$\xi_A = (A_d, B_c)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_c, B_d)^5$$

$$\xi_B = (A_c, B_d)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_d, B_c)^5$$

Prisoners Dilemma: PE, NE

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

PE NE

The paradox of Prisoner's Dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome

Prisoners Dilemma: PE, NE

PE

NE

dominant

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

The paradox of Prisoner's Dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome

Prisoners Dilemma: PE, NE

		B_C	B_D	
social welfare optimal	A_C	1, 1	5, 0	PE
dominant	A_D	0, 5	3, 3	NE

The paradox of Prisoner's Dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome

Mediated Prisoners Dilemma

	Cooperate	Defect
Cooperate	1, 1	5, 0
Defect	0, 5	3, 3

Mediated Game

	Mediator	Cooperate	Defect
Mediator	1, 1	0, 5	2, 2
Cooperate	5, 0	1, 1	5, 0
Defect	2, 2	0, 5	3, 3

Mediated Equlibrium

	Mediator	Cooperate	Defect
Mediator	1, 1	0, 5	2, 2
Cooperate	5, 0	1, 1	5, 0
Defect	2, 2	0, 5	3, 3

Iterated Prisoner Dilemma

- The problem of repeatedly played PD game. Optimization for total count of each player outcome. Sometimes IPD can be played against a range of different; opponents (or even several at the same time).
  - motives for cooperation: (i) if you know you will be meeting your opponent again, then the incentive to defect appears to evaporate. (ii) defection may be punished in the future round,
  - motives for defection: (i) you can test the water by defection (ii) cooperative defection is the rational choice in the infinitely repeated prisoner's dilemma

Iterated Prisoner Dilemma

01

- What strategy to choose, so as to maximize your overall payoff?
- Axelrod (1984) investigated this problem, with a computer tournament for programs playing the iterated prisoner's dilemma:

ALLD	Always defect. the Hawk or Free rider strategy
ALLC	Always cooperate
TITforTAT	first cooperate, than do what your opponent did
TF2T	Same as above, but requires TWO consecutive defections for a defection to be returned
STFT	Suspicious TFT - first, defect. If the opponent retaliated, then play TITforTAT. Otherwise intersperse cooperation & defection.
JOSS	As TIT-FOR-TAT, except periodically defect.

Mixed Strategy

01

- In many games, deterministic strategy is very inefficient.
  - example: matching pennies, security games
- Solution: randomize selection of an action
- Pure strategy
  - agents makes the decision to play one action
- Mixed strategy
  - agents chose to play more actions with positive probabilities
  - support of the mixed strategy is the set of all selected actions
- Payoff
  - given the strategy profile $s \in S$ for all agents, the utility for the agent i

$$u_i(s) = \sum_{a \in A} u_i(a_i) Pr(a|s) \qquad Pr(a|s) = \prod_{j \in N} s_j(a_i)$$

Mixed Strategy

01

Let us generalize the NE concepts for strategy profiles:

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i,\ u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i,\ s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies:

Mixed Strategy

Let us generalize the NE concepts for strategy profiles:

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i$, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i$, $s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies: both players play heads/tails 50%/50%

Mixed Strategy

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the support. For BoS, let's look for an equilibrium where all actions are part of the support.

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play *B* with p, *F* with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

Mixed Strategy

	B	*F*
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- Let player 2 play *B* with p, *F* with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

Mixed Strategy

	B	*F*
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- Let player 2 play *B* with p, *F* with $1-p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

Mixed Strategy

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play *B* with p, *F* with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

$$u_1(B) = u_1(F)$$
$$2p + 0(1-p) = 0p + 1(1-p)$$
$$p = \frac{1}{3}$$

Mixed Strategy

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Likewise: Let player 1 play *B* with q, *F* with $1 - q$.

$$u_2(B) = u_2(F)$$
$$q + 0(1-q) = 0q + 2(1-q)$$
$$q = \frac{2}{3}$$

- Thus the mixed strategy $(\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3})$ is in Nash Equilibrium

Interpreting Mixed Strategy

- What does it mean to play a mixed strategy? Different interpretations:
  - Randomize to confuse your opponent
    - *consider the matching pennies example*
  - Randomize when they are uncertain about the other's action
    - *consider battle of the sexes*
  - Randomize when they you allocated limitted resources
- Mixed strategies are a concise description of what might happen in repeated play: count of pure strategies in the limit
- Mixed strategies describe population dynamics: 2 agents chosen from a population, all having deterministic strategies.