

Multiagent Systems

8. Multiagent Interaction

B. Nebel, C. Becker-Asano, S. Wölfl

Albert-Ludwigs-Universität Freiburg

June 18, 2014

Multiagent Systems

June 18, 2014 — 8. Multiagent Interaction

8.1 Preferences

8.2 Strategic games

8.3 Summary

Multiagent interactions

- ▶ So far: we have looked at agent communication, but not described how it is used in actual agent interactions
- ▶ In itself, communication does not have much effect on the agents
- ▶ Now, we are going to look at interactions in which agents **affect** each other through their actions
- ▶ Assume agents to have “spheres of influence” that they control in the environment
- ▶ Also, we assume that the **welfare** (goal achievement, utility) of each agent at least partially depends on the actions of others
- ▶ This part of the lecture will deal with what agents should **do** in the presence of other agents (which also do stuff)

Example: Bach or Stravinski

A and *E* want to go to a concert tonight. There are two concerts, one in which work by Bach and one in which work by Stravinski is performed.

- ▶ *A* prefers Bach to Stravinski, *E* prefers Stravinski to Bach.
- ▶ Both prefer going together to going alone to a concert.

What are the possible **solutions** in this situation?

8.1 Preferences

Preferences and utilities

We first need an abstract model of interactions:

- ▶ Assume a set of possible outcomes $O = \{o_1, \dots, o_n\}$
(e.g., possible “runs” of the system until final states are reached)

Preference ordering

A **preference ordering** for agent i on O is a binary relation $\succeq_i \subseteq O \times O$ that is reflexive, transitive, and total, i.e.:

- ▶ $o \succeq_i o$
- ▶ $o \succeq_i o'$ and $o' \succeq_i o'' \Rightarrow o \succeq_i o''$
- ▶ for all $o, o' \in O$, either $o \succeq_i o'$ or $o' \succeq_i o$.
- ▶ Such an ordering is used to express the preferences of agent i over O
- ▶ Write $o \succ_i o'$ if $o \succeq_i o'$ and $o' \not\succeq_i o$

Preferences and utilities


- ▶ Preferences are often expressed through a utility function $u_i: O \rightarrow \mathbb{R}$:

$$u_i(o) > u_i(o') \Leftrightarrow o \succ_i o' \text{ and } u_i(o) \geq u_i(o') \Leftrightarrow o \succeq_i o'$$

- ▶ Utilities make representing preferences easier because the ordering follows naturally if we use real numbers.
- ▶ Often, people falsely associate utility directly with money!
- ▶ Intuitively, the utility of money depends on how much money one already has.
- ▶ Therefore, utility does not increase proportionally with monetary wealth.

Utility vs money

- ▶ The utility of money:



- ▶ Empirical evidence suggests utility of money is often very close to logarithm function for humans
- ▶ This shows that utility function depends on agent's risk aversion attitude (value of additional utility depending on current “wealth”)

Multiagent encounters

- ▶ Applying the above to a multiagent setting, we need to consider several agents' actions and the outcomes they lead to
- ▶ For now, restrict ourselves to two players and identical sets of actions
- ▶ Abstract architecture: state transformer function becomes $\tau: Ac \times Ac \rightarrow O$ where Ac is the set of actions available to both agents
- ▶ Outcome depends on other's actions!
- ▶ For pairs $(a_1, a_2), (a'_1, a'_2) \in Ac \times Ac$ we write:

$$(a_1, a_2) \succeq_i (a'_1, a'_2) \text{ if } \tau(a_1, a_2) \succeq_i \tau(a'_1, a'_2)$$

— similarly for \succ , and utilities $u_{1/2}(\tau(a_1, a_2))$

- ▶ We consider agents to be **rational** if they prefer actions that lead to preferred outcomes

8.2 Strategic games

Game theory

- ▶ Game theory: the mathematical theory of interaction problems of this sort
- ▶ Focus: developing solution concepts for games
- ▶ **Basic model**: agents perform simultaneous actions (potentially over several stages), the actual outcome depends on the combination of actions chosen by all agents
- ▶ **Normal-form games**: final result reached in a single step (in contrast to **extensive-form games**)

Game theory

Formal setting:

- ▶ Agents: $\{1, \dots, n\}$
- ▶ Instead of the term "action" we use the term "strategy"
- ▶ S_i : the set of (pure) strategies for agent i ,
- ▶ $S = \prod_{i=1}^n S_i$: the space of **joint strategies** (or: **strategy profiles**)
- ▶ $u_i: S \rightarrow \mathbb{R}$: **utility functions**, maps strategy profiles to utilities
- ▶ **Mixed strategy** of agent i : a probability distribution $\sigma_i: S_i \rightarrow [0, 1]$
- ▶ **Mixed strategy profile** a tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of mixed strategies (one for each agent)

Example: Prisoner's Dilemma

Two men are collectively charged with a crime and held in separate cells, with no way of meeting or communicating. They are told that:

- ▶ if one confesses and the other does not, the confessor will be freed, and the other will be jailed for three years;
- ▶ if both confess, then each will be jailed for two years.

Both prisoners know that if neither confesses, then they will each be jailed for one year.

Payoff matrix for this game:

		Player 2	
		C	D
Player 1	C	2, 2	0, 3
	D	3, 0	1, 1

Dominance and best response strategies

Dominance and best response strategies: two simple and very common criteria for rational decision making in games

Given a strategy profile s , let s_{-i} be the profile $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$. Similar abbrev. used for S .

Dominance

A strategy $s_i \in S_i$ is said to **(strictly) dominate** $s'_i \in S_i$ if for each $s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$.

Notice:

- ▶ Strictly dominated strategies can be safely eliminated from the set of strategies, a rational agent will never play them
- ▶ Some games are solvable in dominant strategy equilibrium, i.e., all agents have a single (pure/mixed) strategy that dominates all other strategies

Example: Iterated elimination

		Player 2		
		A	B	C
Player 1	A	2, 0	1, 1	4, 2
	B	1, 4	1, 1	2, 3
	C	1, 3	0, 2	3, 0

Notice:

- ▶ Result of iterated elimination of strictly dominated strategies does not depend on the elimination order
- ▶ ... does, in general, not hold for weakly dominated strategies

Dominance and Best Response Strategies

Best response

Strategy $s_i \in S_i$ is a **best response** to strategies $s_{-i} \in S_{-i}$ if for each strategy $s'_i \in S_i$, $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$.

- ▶ Weaker notion, only considers optimal reaction to a specific behaviour of other agents
- ▶ Unlike dominant strategies, best-response strategies (trivially) always exist at least one s'_i
- ▶ Replace s_i/s_{-i} above by σ_i/σ_{-i} to extend the definitions for dominant/best-response strategies to mixed strategies

Nash Equilibrium

Nash (1951) defined the most famous equilibrium concept for normal-form games:

Nash equilibrium

A strategy profile $s \in S$ is said to be in **(pure-strategy) Nash equilibrium (NE)** if for each agent $i \in \{1, \dots, n\}$ and each strategy $s'_i \in S_i$, it holds

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}),$$

i.e., no agent has an incentive to deviate from this strategy profile.

Nash Equilibrium

Very appealing notion:

- ▶ It can be shown that a (mixed-strategy) NE always exists

But also some problems:

- ▶ Most of the times: NE is not unique, how to agree on one of them?
- ▶ Proof of existence does not provide method to actually find it
- ▶ Many games do not have pure-strategy NE

Example: Prisoner's Dilemma

In the Prisoner's Dilemma: Nash equilibrium is not Pareto-efficient, i.e., no one will dare to cooperate although mutual cooperation is preferred over mutual defection

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

General conditions on preferences:

- ▶ $DC \succ_1 CC \succ_1 DD \succ_1 CD$

Example: Coordination game

The **Coordination Game**:

		Player 2	
		A	B
Player 1	A	1, 1	-1, -1
	B	-1, -1	1, 1

No temptation to defect, but two equilibria (hard to know which one will be chosen by other party)

The evolution of cooperation?

- ▶ In **zero-sum/constant-sum games** one agent loses what the other wins (e.g., Chess): no potential for cooperation
- ▶ Typical non-zero sum game: there is a potential for cooperation but how should it emerge among self-interested agents?
- ▶ This situation occurs in many real life cases: Nuclear arms race; tragedy of the commons; “Free rider” problems
- ▶ Axelrod’s tournament (1984): a very interesting study of such interaction situations

Axelrod’s tournament

Iterated Prisoner’s Dilemma was played among many different strategies (how to play against different opponents?)

Some of the different strategies:

- ▶ **ALL-D**: always defect, no matter what the other agent has done in the past
- ▶ **RANDOM**: select one of both options with equal probability
- ▶ **TIT-FOR-TAT**: on the first round cooperate; on all later rounds t ; mimic what the other agent has done on round $t - 1$
- ▶ ...

Axelrod’s tournament

- ▶ In single-shot PD, defection is the rational solution
- ▶ In (infinitely) iterated case, cooperation is the rational choice in the PD
- ▶ ... but not if game has a fixed, known length (“backward induction” problem)
- ▶ TIT-FOR-TAT strategy performed best against a variety of strategies (this does not mean it is the best strategy, though!)
- ▶ Axelrod’s conclusions from this: ... (discussed after next exercise sheet)

Example: The Game of Chicken

Some payoff matrix for Chicken

		Player 2	
		C	D
Player 1	C	2, 2	1, 3
	D	3, 1	0, 0

General conditions on preferences:

- ▶ $DC \succ_1 CC \succ_1 CD \succ_1 DD$

8.3 Summary

- Thanks

Critique

While game-theoretic/decision-theoretic approaches are currently very popular, there is also some criticism:

- ▶ How far can we get in terms of cooperation while assuming purely self-interested agents?
- ▶ Good for economic interactions but how about other social processes?
- ▶ In a sense, these approaches assume “worst case” of possible agent behaviour and disregard higher (more fragile) levels of cooperation
- ▶ Although mathematically rigorous,
 - ▶ the proofs only work under simplifying assumptions
 - ▶ often don't consider irrational behaviour
 - ▶ can only deal with a “utilitized” world
- ▶ Relationship to goal-directed, rational reasoning (e.g. BDI) and to deductive reasoning complex and not entirely clear

Summary

- ▶ Discussed simple, abstract models of multiagent encounters
- ▶ Utilities, preferences and outcomes
- ▶ Game-theoretic models and solution concepts
- ▶ Examples: Prisoner's Dilemma, Coordination Game
- ▶ Axelrod's tournament: its conclusions and critique
- ▶ **Next time:** Social Choice

Acknowledgments

These lecture slides are based on the following resources:

- ▶ Dr. Michael Rovatsos, The University of Edinburgh
<http://www.inf.ed.ac.uk/teaching/courses/abs/abs-timetable.html>
- ▶ Michael Wooldridge: **An Introduction to MultiAgent Systems**, John Wiley & Sons, 2nd edition 2009.