

TDDE13 LE1 HT2025

– Introduction to MAS

Fredrik Heintz

**Dept. of Computer Science
Linköping University**


fredrik.heintz@liu.se

@FredrikHeintz


Outline:

- **Multi-agent Systems**
- **Communication**
- **Game Theory**
- **Social Choice**
- **Teamwork**

Multiagent systems



Google AI Scientist



Learning Outcomes

- The overall aim of the course is to give an overview of multiagent systems and in-depth knowledge of some areas of multiagent systems. After the course students should be able to:
 - List and explain important problems and techniques in the area of multiagent systems.
 - Explain how central algorithms in the area of multiagent systems work.
 - Be able to implement some central algorithm in the area of multiagent systems.
 - Evaluate and apply different game theoretic approaches.
 - Design and use auctions for allocating resources in a multiagent system.
 - Model relevant aspects of multiagent system decision making using markov decision processes and logics.

Lectures and Seminars Part I

- The course material for week 2-3 are the following lectures from Game Theory Online:
 - Game Theory I - Week 1 (Introduction)
 - Game Theory I - Week 2 (Mixed-strategy Nash equilibria)
 - Game Theory I - Week 3 (Alternate solution concepts)
 - Game Theory I - Week 4 (Extensive-Form Games)
 - Game Theory I - Week 5 (Repeated Games)
 - Game Theory I - Week 6 (Bayesian Games)
 - Game Theory I - Week 7 (Coalitional Games)
- and the following pre-recorded lectures
 - Cooperative game theory
- *Nov 13, Lab session: Agentic AI*
- *Nov 20, Seminar: Game theory (**exercise set 1**)*
- Dec 15, **Deadline:** Lab 1 report submitted

Lectures and Seminars Part II

- The course material for week 4 and 5 are the following lectures from Game Theory Online:
 - Game Theory II - Week 1 (Social Choice)
 - Game Theory II - Week 2 (Mechanism Design)
 - Game Theory II - Week 3 (VCG)
 - Game Theory II - Week 4 (Auctions)
- and the following pre-recorded lectures
 - Multi-agent learning
- *Dec 4, Lab session: Multi-agent reinforcement learning*
- *Dec 6, Seminar: Social choice, mechanism design and auctions (**exercise set 2**)*
- *Dec 8-9, Discuss individual reports individually with Amath Sow*
- Dec 9, **Deadline:** Approved choice of subject for the individual report
- Dec 15, **Deadline:** Lab 2 report submitted
- *Dec 18, Seminar: Student presentations on individual report*
- Jan 12, **Deadline:** Individual report submitted

Examination

- **LAB 2hp**
 - Agentic AI
 - Multiagent Reinforcement Learning
- **UPG 4hp**
 - **[7 points]** Assignment Set 1: Agents and Game Theory
 - **[7 points]** Assignment Set 2: Mechanism Design, Social Choice, and Coalitional Game Theory
 - 0-7 points grade U; 8-10 points grade 3; 11-12 points grade 4; 13-14 points grade 5
 - **[U,3,4,5]** Individual written report + presentation at seminar
 - The overall grade is the sum divided by two rounded down.

Outline Le 1-2

- Multi-agent Systems
- Communication
- Game Theory
- Social Choice
- Teamwork
- Distributed Problem Solving
- Task Allocation
- Summary

Multi-Agent Systems


“A multi-agent system (MAS) can be defined as a loosely coupled network of problem solvers that interact to solve problems that are beyond the individual capabilities or knowledge of each problem solver.”

Durfee and Lesser, 1989

Characteristics of MAS

- The participants are self-interested.
- The participants and their capabilities changes over time, i.e. open systems.
- Each participant has incomplete information or capabilities for solving the problem and, thus, has a limited viewpoint.
- There is no system global control.
- Data is decentralized.
- Computation is asynchronous.

View of a Canonical MAS



Motivations for MAS

- Model human interactions.
- Solve very large and open problems.
- Interconnect and interoperate legacy system, information sources, or experts.
- Solve problems that naturally can be regarded as a society of agents.
- Enhance computational efficiency, reliability, extensibility, maintainability, flexibility, and reuse.

Multiagent Systems is Interdisciplinary


- The field of Multiagent Systems is influenced and inspired by many other fields:
 - Economics
 - Philosophy
 - Game Theory
 - Logic
 - Ecology
 - Social Sciences
- What makes the multiagent systems field unique is that it emphasizes that the agents in question are *computational, information processing* entities.

Agent Design, Society Design

- **Agent design:** How do we build agents capable of independent, autonomous action, so that they can successfully carry out tasks we delegate to them?
- **Society design:** How do we build agents that are capable of interacting (cooperating, coordinating, negotiating) with other agents in order to successfully carry out those delegated tasks, especially when the other agents cannot be assumed to share the same interests/goals?


Agent-Oriented Programming

	Machine Language	Structured Programming	Object-Oriented Programming	Agent-Oriented Programming
Structural unit	Program	Subroutine	Object	Agent
Relation to previous level		Bound unit of program	Subroutine + persistent local state	Object + independent thread of control + initiative



Intelligent Agents


An **agent** is anything that can be viewed as **perceiving** its **environment** through **sensors** and **acting** upon that environment through **actuators**.



Intelligent Agents

AI Technology areas

- Perception
- Learning
- Knowledge representation and reasoning
- Planning and decision making
- Control



Agent Architectures Summary

- Originally (1956-1985), pretty much all agents designed within AI were *symbolic reasoning* agents
- Its purest expression proposes that agents use *explicit logical reasoning* in order to decide what to do
- Problems with symbolic reasoning led to a reaction against this — the so-called *reactive agents* movement, 1985–present
- From 1990-present, a number of alternatives proposed: *hybrid architectures*, which attempt to combine the best of reasoning and reactive architectures

Deliberative Architectures

Properties


- Internal state (using symbolic representation)
- Search-based decision making
- Goal directed

Benefits

- Nice and clear (logics) semantics
- Easy to analyze by proving properties

Problems

- Can't react in a timely manner to events that requires immediate actions. Intractable algorithms.
- Hard to create a symbolic representation from continuous sensor data. The anchoring problem.



Reactive Agent Architectures

Properties


- No explicit world model
- Rule-based decision making

Benefits

- Efficient
- Robust

Problems

- The local environment must contain enough information to make a decision.
- Easy to build small agents, hard to build agents with many behaviors or rules.
Emergent behavior.



Hybrid Agent Architectures

Properties


- Tries to combine the good parts of both reactive and deliberative architectures.
- Usually layered architectures.

Benefits


- Attacks the problem on different abstraction levels.
- Has the benefits of both architecture types.

Problems


- Hard to combine the different parts.



(a) Horizontal layering




(b) Vertical layering
(One pass control)



(c) Vertical layering
(Two pass control)

HDRC3: A Distributed Hybrid Deliberative/Reactive Architecture for Autonomous Systems




P. Doherty, J. Kvarnström, M. Wzorek, P. Rudol, F. Heintz and G. Conte. 2014.

HDRC3 - A Distributed Hybrid Deliberative/Reactive Architecture for Unmanned Aircraft Systems.
In K. Valavanis, G. Vachtsevanos, editors, Handbook of Unmanned Aerial Vehicles, pages 849–952.

Human and Computational Thinking

Figure 1: A Comparison of System 1 and System 2 Thinking

System 1 "Fast"	System 2 "Slow"
DEFINING CHARACTERISTICS Unconscious Effortless Automatic	DEFINING CHARACTERISTICS Deliberate and conscious Effortful Controlled mental process
WITHOUT self-awareness or control "What you see is all there is."	WITH self-awareness or control Logical and skeptical
ROLE Assesses the situation Delivers updates	ROLE Seeks new/missing information Makes decisions



Research Challenges in MAS

- How to formulate, describe, decompose, and allocate problems and synthesize results among a group of agents.
- How to enable agents to communicate and interact.
- How to make sure that agents act coherently in making decision or taking action.
- How to enable individual agents to represent and reason about the actions, plans, and knowledge of other agents to coordinate with them.
- How to recognize and reconcile disparate viewpoints and conflicting intentions among a collection of agents trying to coordinate their actions.
- How to engineer and constrain practical multi-agent systems.
- How to design technology platforms and development methodologies for MASs.

Communication

Communication

“Communication is the intentional exchange of information brought about by the production and perception of signs drawn from a shared system of conventional signs.”

Russel and Norvig, 1995

Agent Communication

To communicate agents need:

- a common language;
- a common understanding of the knowledge exchanged; and
- the ability to exchange the above information.
- A way of discovering other agents, network addresses and their capabilities.

Speech Acts

- Most treatments of communication in multiagent systems borrow their inspiration from *speech act theory*.
- Speech act theories are *pragmatic* theories of language, i.e., theories of language *use*. They attempt to account for how language is used by people every day to achieve their goals and intentions.
- The origin of speech act theories is usually traced to Austin's 1962 book, *How to Do Things with Words*. He noticed that some utterances are rather like physical actions that appear to change the state of the world.
 - declaring war
 - christening
- A theory of how utterances are used to achieve intentions is a speech act theory.

Types of Speech Act

A speech act consists of a *performative* verb and *propositional content*.

- **Representatives**: such as informing, e.g., “It is raining”
- **Directives**: attempts to get the hearer to do something e.g., “please make the tea”
- **Commisives**: which commit the speaker to doing something, e.g., “I promise to ...”
- **Expressives**: whereby a speaker expresses a mental state, e.g., “thank you!”
- **Declarations**: such as declaring war or christening

Types of Speech Acts

- performative = request
content = "the door is closed"
speech act = "please close the door"
- performative = inform
content = "the door is closed"
speech act = "the door is closed!"
- performative = inquire
content = "the door is closed"
speech act = "is the door closed?"

Plan Based Semantics

- How does one define the semantics of speech acts? When can one say someone has uttered, e.g., a request or an inform?
- Cohen & Perrault (1979) defined semantics of speech acts using the *precondition-delete-add* list formalism of (STRIPS) planning research.
- Note that a speaker can not (generally) *force* a hearer to accept some desired mental state.

Semantics for $Request(A, B, \alpha)$

- Pre-conditions:
 - A believes B can do α
(you don't ask someone to do something unless you think they can do it)
 - A believes B believes B can do α
(you don't ask someone unless they believe they can do it)
 - A believes A wants to do α
(you don't ask someone unless you want it)
- Post-condition:
 - B believes A believes A wants to do α
(the effect is to make them aware of your desire)

Game Theory

Questions that Game Theory tries to answer

- Which action should you use?
- Does it depend on what you think the other person will do?
- What kind of behavior can the designer expect?
- Would any two users behave the same?
- Will this change if users can communicate with each other beforehand?
- Under what changes would the users' decisions still be the same?
- How would you behave if you knew you would face this situation repeatedly with the same person?

Utilities and Preferences

- Assume we have just two agents: $Ag = \{i, j\}$
- Agents are assumed to be *self-interested*:
they *have preferences over how the environment is*
- Assume $\Omega = \{\omega_1, \omega_2, \dots\}$ is the set of “outcomes” that agents have preferences over
- We capture preferences by *utility functions*:


$$u_i = \Omega \rightarrow \mathbb{R}$$

$$u_j = \Omega \rightarrow \mathbb{R}$$

- Utility functions lead to *preference orderings* over outcomes: $\omega \succeq_i \omega'$
means $u_i(\omega) \geq u_i(\omega')$
 $\omega \succ_i \omega'$ means $u_i(\omega) > u_i(\omega')$

What is Utility?

- Utility is *not* money (but it is a useful analogy)
- Typical relationship between utility & money:



Multiagent Encounters

- We need a model of the environment in which these agents will act...
 - agents simultaneously choose an action to perform, and as a result of the actions they select, an outcome in Ω will result
 - the *actual* outcome depends on the *combination* of actions
 - assume each agent has just two possible actions that it can perform, C (“cooperate”) and D (“defect”)
- Environment behavior given by *state transformer function*:

$$\tau : \underbrace{Ac}_{\text{agent } i\text{'s action}} \times \underbrace{Ac}_{\text{agent } j\text{'s action}} \rightarrow \Omega$$

Multiagent Encounters

- Here is a state transformer function:

$$\tau(D, D) = \omega_1 \quad \tau(D, C) = \omega_2 \quad \tau(C, D) = \omega_3 \quad \tau(C, C) = \omega_4$$

(This environment is sensitive to actions of both agents.)

- Here is another:

$$\tau(D, D) = \omega_1 \quad \tau(D, C) = \omega_1 \quad \tau(C, D) = \omega_1 \quad \tau(C, C) = \omega_1$$

(Neither agent has any influence in this environment.)

- And here is another:

$$\tau(D, D) = \omega_1 \quad \tau(D, C) = \omega_2 \quad \tau(C, D) = \omega_1 \quad \tau(C, C) = \omega_2$$

(This environment is controlled by j .)

Rational Action

- Suppose we have the case where *both* agents can influence the outcome, and they have utility functions as follows:

$$\begin{array}{llll} u_i(\omega_1) = 1 & u_i(\omega_2) = 1 & u_i(\omega_3) = 4 & u_i(\omega_4) = 4 \\ u_j(\omega_1) = 1 & u_j(\omega_2) = 4 & u_j(\omega_3) = 1 & u_j(\omega_4) = 4 \end{array}$$

- With a bit of abuse of notation:

$$\begin{array}{llll} u_i(D, D) = 1 & u_i(D, C) = 1 & u_i(C, D) = 4 & u_i(C, C) = 4 \\ u_j(D, D) = 1 & u_j(D, C) = 4 & u_j(C, D) = 1 & u_j(C, C) = 4 \end{array}$$

- Then agent i 's preferences are:
- “ C ” $\vdash C, C \succeq_i C, D \succ_i D, C \succeq_i D, D$ cause i prefers all outcomes that arise through C over all outcomes that arise through D .)

Payoff Matrices

- We can characterize the previous scenario in a *payoff matrix*:

		i	
		defect	coop
j	defect	1 1	4 1
	coop	1 4	4 4

- Agent i is the *column player*
- Agent j is the *row player*
- Often called *Normal Form*.

Solution Concepts

- How will a rational agent behave in any given scenario?
- Answered in *solution concepts*:
 - dominant strategy;
 - Nash equilibrium strategy;
 - Pareto optimal strategies;
 - strategies that maximize social welfare.

Dominant Strategies

- Given any particular strategy (either C or D) of agent i , there will be a number of possible outcomes
- We say s_1 *dominates* s_2 if every outcome possible by i playing s_1 is preferred over every outcome possible by i playing s_2
- A rational agent will never play a dominated strategy
- So in deciding what to do, we can *delete dominated strategies*
- Unfortunately, there isn't always a unique undominated strategy

		i	
		defect	coop
j	defect	1 1	4 4
	coop	1 4	4 4

		i	
		defect	coop
j	defect	1 1	0 0
	coop	0 0	1 1

Pareto Optimality

- An outcome is said to be *Pareto optimal* (or *Pareto efficient*) if *there is no other outcome that makes one agent better off without making another agent worse off*.
- If an outcome is Pareto optimal, then at least one agent will be reluctant to move away from it (because this agent will be worse off).
- If an outcome ω is *not Pareto optimal*, then there is another outcome ω' that *makes everyone as happy, if not happier*, than ω .
- “Reasonable” agents would agree to move to ω' in this case. (Even if I don't directly benefit from ω , you can benefit without me suffering.)

(Pure Strategy) Nash Equilibrium

- In general, we will say that two strategies s_1 and s_2 are in Nash equilibrium if:
 1. under the assumption that agent i plays s_1 , agent j can do no better than play s_2 ; and
 2. under the assumption that agent j plays s_2 , agent i can do no better than play s_1 .
- *Neither agent has any incentive to deviate from a Nash equilibrium*
- Unfortunately:
 1. *Not every interaction scenario has a Nash equilibrium*
 2. *Some interaction scenarios have more than one Nash equilibrium*

		i	
		defect	coop
j	defect	1	4
	coop	1	4

		i	
		defect	coop
j	defect	1	0
	coop	0	1

Social Welfare

- The social welfare of an outcome ω is the sum of the utilities that each agent gets from ω :

$$\sum_{i \in Ag} u_i(\omega)$$

- Think of it as the “**total amount of utility in the system**”.
- As a solution concept, may be appropriate when the whole system (all agents) has a single owner (then overall benefit of the system is important, not individuals).

The 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



		i	
		defect	coop
j	defect	2 1	2 4
	coop	4 1	3 3



The Prisoner's Dilemma

- Solution concepts
 - D is a dominant strategy.
 - (D,D) is the only Nash equilibrium.
 - All outcomes except (D,D) are Pareto optimal.
 - (C,C) maximizes social welfare.

		i	
		defect	coop
j	defect	2 2	1 4
	coop	4 1	3 3

- The *individual rational* action is *defect*

This guarantees a payoff of no worse than 2, whereas cooperating guarantees a payoff of at most 1. So defection is the best response to all possible strategies: both agents defect, and get payoff = 2

- But *intuition* says this is *not* the best outcome:
Surely they should both cooperate and each get payoff of 3!

The Prisoner's Dilemma

- This apparent paradox is *the fundamental problem of multi-agent interactions*.
It appears to imply that *cooperation will not occur in societies of self-interested agents*.
- Real world examples:
 - nuclear arms reduction (“why don’t I keep mine. . .”)
 - free rider systems — public transport;
 - television licenses.
- The prisoner’s dilemma is *ubiquitous*.
- Can we recover cooperation?

The Iterated Prisoner's Dilemma

- One answer: *play the game more than once*
- If you know you will be meeting your opponent again, then the incentive to defect appears to evaporate
- *Cooperation is the rational choice in the infinitely repeated prisoner's dilemma*

Backwards Induction

- But..., suppose you both know that you will play the game exactly n times
 - On round $n-1$, you have an incentive to defect, to gain that extra bit of payoff...
 - But this makes round $n-2$ the last “real”, and so you have an incentive to defect there, too.
 - This is the *backwards induction* problem.
- Playing the prisoner’s dilemma with a fixed, finite, pre-determined, commonly known number of rounds, defection is the best strategy

Axelrod's Tournament

- Suppose you play iterated prisoner's dilemma against a *range* of opponents...
What strategy should you choose, so as to maximize your overall payoff?
- Axelrod (1984) investigated this problem, with a computer tournament for programs playing the prisoner's dilemma

Strategies in Axelrod's Tournament

- ALLD:
 - “Always defect” — the *hawk* strategy;
- TIT-FOR-TAT:
 1. On round $u = 0$, cooperate
 2. On round $u > 0$, do what your opponent did on round $u-1$
- TESTER:
 - On 1st round, defect. If the opponent retaliated, then play TIT-FOR-TAT. Otherwise intersperse cooperation and defection.
- JOSS:
 - As TIT-FOR-TAT, except periodically defect

Recipes for Success in Axelrod's Tournament

- Axelrod suggests the following rules for succeeding in his tournament:
 - *Don't be envious:*
Don't play as if it were zero sum!
 - *Be nice:*
Start by cooperating, and reciprocate cooperation
 - *Retaliate appropriately:*
Always punish defection immediately, but use “measured” force — don't overdo it
 - *Don't hold grudges:*
Always reciprocate cooperation immediately


Competitive and Zero-Sum Interactions


- Where preferences of agents are diametrically opposed we have *strictly competitive* scenarios
- Zero-sum encounters are those where utilities sum to zero:
$$u_i(\omega) + u_j(\omega) = 0 \quad \text{for all } \omega \in \Omega$$
- Zero sum implies strictly competitive
- Zero sum encounters in real life are very rare, but people tend to act in many scenarios as if they were zero sum

Matching Pennies

- Players i and j simultaneously choose the face of a coin, either “heads” or “tails”.
- If they show the same face, then i wins, while if they show different faces, then j wins.

	i heads	i tails
j heads	1 -1	-1 1
j tails	-1 1	1 -1





	B: Rock	B: Scissors	B: Paper
A: Rock	A: 0; B: 0	A:1; B: -1	A: -1; B: 1
A: Scissors	A:-1; B: 1	A: 0; B: 0	A: 1; B: -1
A: Paper	A: 1; B: -1	A: -1; B: 1	A: 0; B: 0

Mixed Strategies for Matching Pennies

- No pair of strategies forms a pure strategy Nash Equilibrium: whatever pair of strategies is chosen, somebody will wish they had done something else.
- The solution is to allow *mixed strategies*:
 - play “heads” with probability 0.5
 - play “tails” with probability 0.5.
- This is a Nash Equilibrium strategy.

Mixed Strategies

- A mixed strategy has the form
 - play α_1 with probability p_1
 - play α_2 with probability p_2
 - ...
 - play α_k with probability p_k .

such that $p_1 + p_2 + \dots + p_k = 1$.

- Nash proved that *every finite game has a Nash equilibrium in mixed strategies*.

Computing the Optimal Strategy

- Let $G=(\{1,2\}, A_1 \times A_2, (u_1, u_2))$ be a zero sum game and U^*_i the expected reward for player i when the game is in an equilibrium.
- Then it follows that $U^*_1 = - U^*_2$.

$$\begin{array}{ll} \text{minimize} & U_1^* \\ \text{subject to} & \sum_{k \in A_2} u_1(a_1^j, a_2^k) \cdot s_2^k \leq U_1^* \quad \forall j \in A_1 \\ & \sum_{k \in A_2} s_2^k = 1 \\ & s_2^k \geq 0 \quad \forall k \in A_2 \end{array}$$

Social Choice

Social Choice

- *Social choice theory is concerned with group decision making.*
- Classic example of social choice theory: *voting*.
- Formally, the issue is *combining preferences to derive a social outcome*.

Components of a Social Choice Model

- Assume a set $Ag = \{1, \dots, n\}$ of *voters*.
- These are the entities who expresses preferences.
- Voters make group decisions wrt a set $\Omega = \{\omega_1, \omega_2, \dots\}$ of *outcomes*.
Think of these as the *candidates*.
- If $|\Omega| = 2$, we have a *pair wise election*.

Preferences

- Each voter has preferences over Ω : an *ordering over* the set of possible outcomes Ω .
- Example. Suppose
$$\Omega = \{gin, rum, brandy, whisky\}$$
then we might have agent F with *preference order*:
$$\omega_F = (brandy, rum, gin, whisky)$$
meaning
$$brandy >_F rum >_F gin >_F whisky$$

Preference Aggregation

- The fundamental problem of social choice theory:
Given a collection of preference orders, one for each voter, how do we combine these to derive a group decision, that reflects as closely as possible the preferences of voters?
- Two variants of preference aggregation:
 - *social welfare functions;*
 - *social choice functions.*

Social Welfare Functions

- Let $\Pi(\Omega)$ be the set of preference orderings over Ω .
- A *social welfare function* takes the *voter preferences* and produces a *social preference order*:

$$f : \underbrace{\Pi(\Omega) \times \cdots \times \Pi(\Omega)}_{n \text{ times}} \rightarrow \Pi(\Omega).$$

- We let $\succ^*$ denote the outcome of a social welfare function
 - Example: beauty contest.

Social Choice Functions

- Sometimes, we just one to select *one of the possible* candidates, rather than a social order.
- This gives *social choice functions*:

$$f : \underbrace{\Pi(\Omega) \times \cdots \times \Pi(\Omega)}_{n \text{ times}} \rightarrow \Omega.$$

- Example: presidential election.

Voting Procedures: Plurality

- Social choice function: selects a single outcome.
- Each voter submits preferences.
- Each candidate gets one point for every preference order that ranks them first.
- Winner is the one with largest number of points.
 - Example: Political elections in UK.
- If we have only two candidates, then plurality is a *simple majority election*.

Anomalies with Plurality

- Suppose $|Ag| = 100$ and $\Omega = \{\omega_1, \omega_2, \omega_3\}$ with:
 - 40% voters voting for ω_1
 - 30% of voters voting for ω_2
 - 30% of voters voting for ω_3
- With plurality, ω_1 gets elected even though a *clear majority (60%) prefer another candidate!*

Strategic Manipulation by Tactical Voting

- Suppose your preferences are

$$\omega_1 \succ_i \omega_2 \succ_i \omega_3$$

while you believe 49% of voters have preferences

$$\omega_2 \succ_i \omega_1 \succ_i \omega_3$$

and you believe 49% have preference

$$\omega_3 \succ_i \omega_2 \succ_i \omega_1$$

You may do better voting for ω_2 , *even though this is not your true preference profile*.

- This is *tactical voting*: an example of strategic manipulation of the vote.

Condorcet's Paradox

- Suppose $Ag = \{1, 2, 3\}$ and $\Omega = \{\omega_1, \omega_2, \omega_3\}$ with:

$$\omega_1 \succ_1 \omega_2 \succ_1 \omega_3$$

$$\omega_3 \succ_2 \omega_1 \succ_2 \omega_2$$

$$\omega_2 \succ_3 \omega_3 \succ_3 \omega_1$$

- For every possible candidate, there is another candidate that is preferred by a majority of voters!
- This is *Condorcet's paradox*: there are situations in which, *no matter which outcome we choose, a majority of voters will be unhappy with the outcome chosen*.

Coalitional Games

- *Coalitional games model scenarios where agents can benefit by cooperating.*
- Issues in coalitional games (Sandholm et al, 1999):
 - Coalition structure generation.
 - Teamwork.
 - Dividing the benefits of cooperation.

Teamwork

Working Together

- Why and how should agents work together?
 - Since agents are autonomous, they have to make decisions at *run-time* and be capable of *dynamic coordination*.
- Cooperation is *working together as a team to achieve a shared goal*.
 - Often prompted either by the fact that no one agent can achieve the goal alone, or that cooperation will obtain a better result (e.g., get result faster).
- Coordination is *managing the interdependencies between activities*.
- Negotiation is *the ability to reach agreements on matters of common interest*.
 - Typically involves offer and counter-offer, with compromises made by participants.
- Overall they will need to be able to share:
 - Tasks
 - Information
- If agents are designed by different individuals, they may not have common goals.


Criteria for Assessing Agent-based Systems

- *Coherence* - how well the [multiagent] system behaves as a unit along some dimension of evaluation (Bond and Gasser).
 - We can measure coherence in terms of solution quality, how efficiently resources are used, conceptual clarity and so on.
- *Coordination* - the degree... to which [the agents]... can avoid “extraneous” activity [such as]... synchronizing and aligning their activities (Bond and Gasser).
 - If the system is perfectly coordinated, agents will not get in each others’ way, in a physical or a metaphorical sense.


Cooperative Problem Solving

- How does a group of agents work together to solve problems?
- If we “own” the whole system, we can design agents to help each other whenever asked. In this case, we can assume agents are benevolent: our best interest is their best interest.
- Problem-solving in benevolent systems is *cooperative distributed problem solving* (CDPS).
- There are three stages:
 - Problem decomposition
 - Sub-problem solution
 - Answer synthesis


Task Sharing and Result Sharing



decomposition




solution




synthesis

Task Sharing and Result Sharing



Task sharing



Result sharing


Handling Inconsistency

- A group of agents may have inconsistencies in their:
 - Beliefs, goals or intentions
- Inconsistent beliefs arise because agents have different views of the world.
 - May be due to sensor faults or noise or just because they can't see everything.
- Inconsistent goals may arise because agents are built by different people with different objectives.
- Three ways to handle inconsistency (Durfee et al.)
 - Do not allow it to occur.
 - Build systems that degrade gracefully in the presence of inconsistency.
 - Resolve inconsistencies through negotiation.

Coordination

- Von Martial suggested that *positive coordination* is:
 - Requested (explicit)
 - Non-requested (implicit)
- Non-requested coordination relationships can be as follows.
 - Action equality: we both plan to do something, and by recognizing this one of us can be saved the effort.
 - Consequence: What I plan to do will have the side-effect of achieving something you want to do.
 - Favor: What I plan to do will make it easier for you to do what you want to do.

Coordination Relationships



Coordination

- Partial global planning
- Joint intentions
- Mutual modeling
- Norms and social laws

Allocating Scarce Resources

- Allocation of *scarce resources* amongst a number of agents is central to multiagent systems.
- Resource might be:
 - a physical object
 - the right to use land
 - computational resources (processor, memory, . . .)
- If the resource isn't scarce, there is no trouble allocating it.
- If there is no competition for the resource, then there is no trouble allocating it.

What is an Auction?

- Concerned with *traders* and their *allocations* of:
 - Units of an indivisible *good*; and
 - Money, which is divisible.
- Assume some initial allocation.
- *Exchange* is the free alteration of allocations of goods and money between traders

Limit Price

- Each trader has a value or *limit price* that they place on the good.
 - A buyer who exchanges more than their limit price for a good makes a loss.
 - A seller who exchanges a good for less than their limit price makes a loss.
- Limit prices clearly have an effect on the behavior of traders.
- There are several models, embodying different assumptions about the nature of the good.

Limit Price

- **Private value**

- Good has an value to me that is independent of what it is worth to you.
- Textbook gives the example of John Lennon's last dollar bill.

- **Common value**

- The good has the same value to all of us, but we have differing estimates of what it is.
- Winner's curse

- **Correlated value**

- Our values are related.
- The more you are prepared to pay, the more I should be prepared to pay.

Auctions

- A *market institution* defines how the exchange takes place.
- The change of allocation is market *clearing*.
- Difference between allocations is *net trade*.
 - Component for each trader in the market.
 - Each trader with a non-zero component has a *trade* or *transaction* price.
 - Absolute value of the money component divided by the good component.
- Traders with positive good component are *buyers*
- Traders with negative good component are *sellers*
- *One way traders* are either buyers or sellers but not both.

Yes, but what is an auction?

- An *auction* is a market institution in which messages from traders include some price information – this information may be an offer to buy at a given price, in the case of a *bid*, or an offer to sell at a given price, in the case of an *ask* – and which gives priority to higher bids and lower asks.
- This definition, as with all this terminology, comes from Dan Friedman.

Single versus Multi-dimensional

- Single dimensional auctions
 - The only content of an offer are the price and quantity of some specific type of good.
 - “I’ll bid \$200 for those 2 chairs”
- Multi dimensional auctions
 - Offers can relate to many different aspects of many different goods.
 - “I’m prepared to pay \$200 for those two red chairs, but \$300 if you can deliver them tomorrow.”

Single versus Double-sided

- Single-sided markets
 - Either one buyer and many sellers, or one seller and many buyers.
 - The latter is the thing we normally think of as an auction.
- Two-sided markets
 - Many buyers and many sellers.
- Single sided markets with one seller and many buyers are “sell-side” markets.
- Single-sided markets with one buyer and many sellers are “buy-side”.

Open-cry versus Sealed-bid

- Open cry
 - Traders announce their offers to all traders
- Sealed bid
 - Only the auctioneer sees the offers.
- Clearly as a bidder in an open-cry auction you have more information.
- In some auction forms you pay for preferential access to information.

Single-unit versus Multi-unit

- How many units of the same good are we allowed to bid for?
- Single unit
 - One at a time.
 - Might repeat if many units to be sold.
- Multi-unit
 - Bid both price and quantity.
- “Unit” refers to the indivisible unit that we are selling.
 - Single fish versus box of fish.

First price versus k th price

- Does the winner pay the highest price bid, the second highest price, or the k th highest price?

Single item versus Multi-item

- Not so much quantity as heterogeneity.
- Single item
 - Just the one indivisible thing that is being auctioned.
- Multi-item
 - Bid for a bundle of goods.
 - “Two red chairs and an orange couch, or a purple beanbag.”
 - Valuations for bundles are not linear combinations of the values of the constituents.

Standard Auction Types


- English auction
- Dutch auction
- First-price sealed bid auction
- Vickrey auction

Combinatorial Auctions

- Auctions for bundles of goods.
- A good example of bundles of good are spectrum licenses.
 - For the 1.7 to 1.72 GHz band for Brooklyn to be useful, you need a license for Manhattan, Queens, and Staten Island.
 - Most valuable are the licenses for the same bandwidth.
 - But a different bandwidth license is more valuable than no license

Summary

- A multi-agent system (MAS) can be defined as a loosely coupled network of problem solvers that interact to solve problems that are beyond the individual capabilities or knowledge of each problem solver.
- Communication
- Game Theory
- Social Choice
- Teamwork
- Task Allocation



TDDE13 MAS LE1 HT2025:
Multi-agent Systems
Communication
Game Theory
Social Choice
Teamwork

www.ida.liu.se/~TDDE13