

Multiagent Systems

1. Introduction

B. Nebel, C. Becker-Asano, S. Wölfl

Albert-Ludwigs-Universität Freiburg

2nd of May 2014

Multiagent Systems

2nd of May 2014 — 1. Introduction

1.1 Intro to MAS

Intro to MAS

1.1 Intro to MAS

- What is MAS about?
- What is an agent?
- Agent Autonomy
- Environments
- Intelligent agents
- Delineation
- Summary
- Thanks
- Literature

Intro to MAS What is MAS about?

Multiagent Systems?

Agent

An **agent** is anything that is capable of performing independent actions ...
(... we will need to revisit this “definition”)

MultiAgent System (MAS)

A **multiagent system** is a system that consists of a number of **agents**, which **interact** with one-another.

To successfully interact, agents will require the ability to **cooperate**, **coordinate**, and **negotiate** with each other

Five pervasive trends in computing history

1. Ubiquity:

Cost of processing power decreases dramatically (e.g. Moore's Law), computers used everywhere

2. Interconnection:

Formerly only user-computer interaction, nowadays distributed/networked systems (Internet etc.)

3. Complexity:

Elaboration of tasks carried out by computers has grown

4. Delegation:

Giving control to computers even in safety-critical tasks (aircraft/nuclear plant control)

5. Human-orientation:

Increasing use of metaphors that better reflect human intuition from everyday life (e.g. GUIs, speech recognition, object orientation)

Current challenges for computer systems

Traditional design problem

How to design a system that produces correct output given some input?

Modern-day design problem

How to build a system that:

- ▶ operates independently on someone's behalf
- ▶ in a networked, distributed, large-scale environment
- ▶ in which it interacts with different components pertaining to other users?

One current research focus

Distributed systems in which different components have different goals and need to cooperate

Multiagent Systems

Two fundamental ideas:

1. Individual agents are capable of **autonomous actions** to a certain extent: agents need not to be told exactly what to do
2. **Agents interact** with each other in multiagent systems, which may represent users with different goals

Foundational problems in MAS research:

- ▶ The **agent design problem**: how should agents act to carry out their tasks?
- ▶ The **society design problem**: how should agents interact to carry out their tasks?

These are the **micro** and **macro** perspective of MAS, respectively.

A pure engineering task?

Artificial Intelligence (AI) and MAS have a goal in common

To understand how societies of intelligent beings work

Questions related to this:

- ▶ How can **cooperation** emerge among self-interested agents?
- ▶ How can agents **coordinate** activities with others?
- ▶ What **languages** should agents use to exchange information necessary to organize interaction in a meaningful way?
- ▶ How should they resolve their **conflicts**?
- ▶ How to detect and deal with agents violating **social rules**?

Philosophically speaking

MAS research departs from traditional engineering, because **control** is replaced by **communication**

Some applications of multiagent systems

Two major areas of application:

- ▶ Distributed systems (agents as processing nodes)
- ▶ Personal software assistants (aiding the user)

A variety of subareas:

- ▶ Workflow/business process management
- ▶ Distributed sensing
- ▶ Information retrieval and management
- ▶ Electronic commerce
- ▶ Human-computer interfaces
- ▶ Virtual environments
- ▶ Social simulation
- ▶ ...

What is an agent?

Definition 1 (Russell & Norvig, p. 34)

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




Figure: An agents interacts with an environment through sensors and actuators (after Russel & Norvig, p. 35)

What is an agent?

Definition 1 (Russell & Norvig, p. 34)

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

- ▶ Most widely accepted definition
- ▶ Focus on **situatedness** in the environment
⇒ **embodiment**
- ▶ The agent can only **influence** the environment, but **not fully control** it
⇒ sensor/effector failure, non-determinism

What is an agent?

Definition 2 (Wooldridge, p. 21)

An agent is a computer system that is **situated** in some **environment**, and that is capable of **autonomous action** in this environment in order to meet its **design objectives**

- ▶ Adds a second dimension to the agent definition, i.e. the **relationship between agent and designer/user**
 - ▶ Agent is capable of independent action
 - ▶ Agent action is purposeful
- ▶ Broad consensus that **autonomy** is a central, distinguishing property of agents
⇒ But, this property is also the **most disputed** one ...

Is a thermostat autonomous (enough)?




Figure: The Nest Thermostat (picture taken from <https://nest.com/thermostat/life-with-nest-thermostat/>)

From their website

“Nest learned from the efficient temperatures Johnny and Kate set—like 66 °F/19 °C at night in winter—and programmed itself to help them save. Before Nest, Johnny and Kate just left the house at 70 °F/21 °C year round.”

Would we call it an agent?

Agent Autonomy

Autonomy is a prerequisite for

- ▶ delegating complex tasks to agents
- ▶ ensuring flexible actions in unpredictable environments

(Different definitions highlight different aspects)

Autonomy dilemma

How to make an agent smart without losing control over it!

⇒ **Adjustable Autonomy** (Wooldridge, p. 23)

Transfer control back to user, whenever:

- ▶ agent believes that human's decision will gain **substantially higher benefit**
- ▶ high degree of **uncertainty present**
- ▶ autonomous decision might **cause harm**
- ▶ agent **lacks** decision making **capability**

A practical example on autonomy

NASA's Deep Space 1/2 space probe:

When a space probe makes its long flight from Earth to the outer planets, a ground crew is usually required to continually track its progress, and decide how to deal with unexpected eventualities. This is costly and, if decisions are required quickly, it is simply not practicable. For these reasons, organizations like NASA are seriously investigating the possibility of **making probes more autonomous** – giving them richer decision making capabilities and responsibilities. (Wooldridge)

Properties of agent environments

Environments can be classified according to the following properties (cf. Russel & Norvig, pp. 42–44):

1. Fully observable vs. partially observable
2. Single agent vs. multi agent
3. Deterministic vs. stochastic
4. Episodic vs. sequential
5. Static vs. dynamic
6. Discrete vs. continuous

Properties of agent environments/examples

Examples of task environments and their characteristics:

Task Environment	Observ.	Agents	Determin.	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Determin.	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Determin.	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Backgammon	Fully	Multi	Stochastic	Sequential	Static	Discrete
Taxi driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Determin.	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Interact. tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

(Russell & Norvig, p. 45)

Intelligent agents

The previous definitions state some basic properties of agents, but don't say anything about **intelligent** agents

⇒ No general definition of intelligent agency is needed, but practical criteria that matter in target application scenarios

Not easy to answer, but **desirable properties** include (Wooldridge, pp. 26–27):

- ▶ **Reactivity**: intelligent agents should respond in a timely fashion to changes they perceive in their environment
- ▶ **Proactiveness**: intelligent agents can take the initiative to meet their design objectives, and they exhibit goal-directed behavior
- ▶ **Social ability**: intelligent agents can interact with other agents (and humans) to satisfy their design objectives

Proactiveness + reactivity

Example: The dung beetle

After digging its nest and laying its eggs, it fetches a ball of dung from a nearby heap to plug the entrance; if the ball of dung is removed from its grasp *en route*, the beetle continues on and pantomimes plugging the nest with the nonexistent dung ball, never noticing that it is missing (quoted from Russell & Norvig)

- ▶ Truly flexible behavior hard to achieve
- ▶ Trade-off between the two aspects, because:
 - ▶ Environments are not fixed
⇒ must be able to react to changes, which involves monitoring own activity and environment, etc.
 - ▶ Need for goal-oriented, planned activity
⇒ not sufficient to respond to current circumstances

Social ability

- ▶ Most real-world applications inhabited by multiple agents
- ▶ Single agents have limited resources/capabilities
⇒ goal achievement might require others' (in)action
- ▶ Social ability: the ability to manage one's interactions effectively, as compared to simple exchange of messages (like IPC, or RPC) between processes
- ▶ Shared goals and cooperation:

In the human world, comparatively few of our meaningful goals can be achieved without the *cooperation* of other people, who cannot be assumed to *share* our goals. [...] To achieve our goals in such situations, we must *negotiate* and *cooperate* with others. (Wooldridge, p. 28)

What is agent technology?

Agents as a software engineering paradigm

- ▶ **Interaction** primary aspect of complex software systems
- ▶ Ideal for **loosely coupled** “black-box” components

Agents as a tool for understanding human societies

- ▶ Human society is very complex, computer simulation can be useful
- ▶ Has given rise to (agent-based) **social simulation**

Agents vs. distributed systems

- ▶ Long tradition of distributed systems research
- ▶ But MAS are not simply distributed systems, because of **different goals**

Agents vs. economics/game theory

- ▶ **Distributed rational decision making** extensively studied in economics, game theory very popular
- ▶ Many strengths but also objections

Agents vs. AI

- ▶ Agents grew out of “distributed” AI
- ▶ Much debate whether MAS is sub-field of AI or vice versa
- ▶ **AI** mostly concerned with the **building blocks** of intelligence
⇒ **reasoning** and problem-solving, **planning** and learning, perception and action
- ▶ **Agents** field more concerned with:
 - ▶ **Combining these concepts** (this may mean we have to solve all problems of AI, but agents can also be built without any AI)
 - ▶ **Social interaction**, which has mostly been ignored by standard AI (and is an important part of human intelligence)
- ▶ Agents are a lot about integration (of abilities in one agent or of agents in one environment)

How are agents different?

Agents vs. objects

- ▶ Objects exhibit control over their state but not over their behavior (limited sense of autonomy)
- ▶ Objects do it for free, agents do it because they want to (or for money)

Agents vs. expert systems

- ▶ Expert systems are knowledge-based systems capable of problem solving in rich, complex domains
- ▶ They can be intelligent, but they lack situatedness and usually don't cooperate with each other

Agents vs. intentional systems

- ▶ “The intentional stance”: ascribing mental attitudes to machines (beliefs, intentions, goals etc)
- ▶ Not necessary when simpler model is available but may aid human understanding

Summary

Topics discussed so far:

- ▶ Trends in computing, new challenges
- ▶ Fundamental issues of MAS research
- ▶ Relationships to other fields of research
- ▶ Autonomy: a difficult notion
- ▶ Environments for agents
- ▶ Advanced properties of intelligent agents
- ▶ What agents are and are not

⇒ Next time: Abstract Agent Architectures

Acknowledgments

The lecture slides are partly based on slides developed by Dr. **Michael Rovatsos**, The University of Edinburgh, for his lecture on “Agent-based systems” used with permission.

Literature I



Wooldridge,
An Introduction to MultiAgent Systems,
Wiley & Sons , 2nd edition, 2009.



Russell & Norvig,
Artificial Intelligence: A Modern Approach,
Prentice Hall, 3rd edition, 2010.