Multiagent Systems: Spring 2006

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

The MAS Course in a Nutshell

- A multiagent system (MAS) is a system consisting of several autonomous entities, called *agents*, that interact with each other to either further their own interests (competition) or in pursuit of a joint goal (cooperation).
- While classical Artificial Intelligence has concentrated on modelling (specific aspects of) single agents, the field of multiagent systems focusses on the *interaction* between different agents.
- This course will exemplify some of the core contributions to the theory of multiagent systems made by different disciplines, including *logic*, *economics* and *computer science*.
- Topics to be addressed include: negotiation, multiagent resource allocation, fair division, combinatorial auctions, mechanism design, and preference representation in combinatorial domains, . . .

Organisational Matters

- **Timetable:** Mondays 11am-1pm in I.001 with occasional additional meetings on Fridays 3-5pm in P.016
- **Examination:** There will be several *coursework assignments* on the material covered in the course. In the second block, each student will have to study a *recent paper* from the MAS literature, write a short paper about it, and present their findings in a talk.
- **Website:** Lecture slides, coursework assignments, and other important information will be posted on the course website:

  `http://www.illc.uva.nl/~ulle/teaching/mas/`

  Or find it yourself: `ILLC > People > Ulle > Teaching > ...`

Plan for Today

- **Part I:** Broad overview of the MAS research area as a whole
- **Part II:** Introduction to the issues we will address in this course

Part I

The MAS Research Area

Research in "Distributed AI" started over 25 years ago, but only in the mid 1990s has it become a (or even *the*) major research trend in AI.

Now the main conference (AAMAS) attracts around 550 submissions (of which 20-25% get accepted) and around 750 participants each year. In addition, there are dozens of smaller workshops and conferences.

So, it's a *large*, *young* and *dynamic* research community. That means:

+ Comparatively easy to get into the field.
+ People are open to new ideas and welcome interdisciplinary work.
– Not always easy to see what's good research and what isn't.
– You never know how long it will last ...

What is an agent?

In fact, there's *no definition* that would be commonly agreed upon. But following Wooldridge and Jennings (1995), we can at least say that an agent is a computer system that is:

- *autonomous* — it has control over its own actions
- *reactive* — it reacts to events in its environment
- *proactive* — it acts on its own initiative
- *social* — it interacts with other agents

M. Wooldridge and N.R. Jennings. *Intelligent Agents: Theory and Practice*. Knowledge Engineering Review, 10(2):115–152, 1995.

BDI Architectures

A common approach is to specify agents in terms of:

- a set of *beliefs* about the world,
- a set of *desires* (or goals), and
- a set of *intentions*.

There's also been a lot of work on *axiomatising* the relationships between beliefs, desires, and intentions in many-dimensional modal logics.

A.S. Rao and M.P. Georgeff. *An Abstract Architecture for Rational Agents*. Proc. KR-1992.

Agent Communication

- Communication is a central issue in multiagent systems.
- There have been many different proposals for agent communication languages (ACLs), but a *message* would typically have at least the following components:
  - a *performative* such as `inform`, `request`, or `accept`
  - the actual *content* of the message (application-dependent)
  - names of *sender* and *receivers*, maybe a *timestamp*, ...
- *Semantics:* Early work in particular has tried to explain the meaning of ACLs in terms of mentalistic notions. Examples:
  - Ann may send an `inform`(X)-message to Bob only if she herself believes X to be true.
  - Bob may send a `request`(Y)-message to Ann only if he believes that Ann does not already intend to perform action Y.

Agent Communication (cont.)

- Two schools of thought: *"mentalistic"* vs. *"conventionalist"* approach to agent communication
- *Mental* attitudes (beliefs, intentions) are useful to explain *why* agents may behave in certain ways, but (being non-verifiable) they cannot serve as a basis for building open systems that allow for meaningful communication.
- A somewhat more promising approach to agent communication relies on public norms and *conventions* as a means of specifying the rules of social interaction.

M.P. Singh. *Agent Communication Languages: Rethinking the Principles.* IEEE Computer, 31(12):40–47, 1998.

Interaction Protocols

In the convention-based approach, *protocols* specify the range of *legal follow-ups* available to the participating agents in a given dialogue.

Example for a protocol based on a finite state machine:

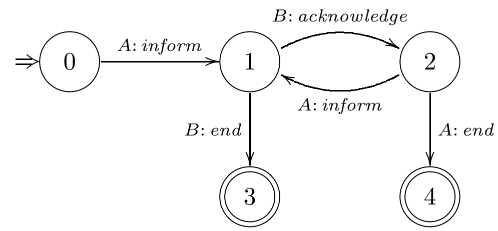

By referring to a protocol (rather than to mental states) we can give a *"social" semantics* to the interactions taking place in a MAS.

In open systems, *public* protocols and agent's *private* strategies may not always match ($\leadsto$ *conformance checking*).

Distributed Problem Solving

Imagine a multiagent system inhabited by several agents with different problem solving capabilities. A particular complex problem may not be solvable by any single agent, but possibly by several agents together.

- *Problem decomposition:* Decompose the original problem into smaller subproblems, that can each be handled by a single agent.
- *Solving each subproblem:* Each agent solves the problems assigned to them. Agents may share information during this stage.
- *Solution synthesis:* Integrate the solutions to the subproblems to arrive at a solution of the overall problem.

In this strand of work people generally assume that agents are *cooperative* and *benevolent* ...

R. Davies and R.G. Smith. *Negotiation as a Metaphor for Distributed Problem Solving.* Artificial Intelligence, 20(1):63–109, 1983.

Negotiation

- Also when agents are *self-interested*, they need to be able to coordinate their actions, resolve conflicts, reach agreements ... they need to be able to *negotiate*.
- There's been much work on *agent-mediated electronic commerce*, but the general techniques actually have much wider appeal.
- Negotiation will be one of the central topics covered in this course.

Agent-oriented Software Engineering

This is a subarea of both MAS research and software engineering, where people are trying to develop formal as well as practical approaches to building distributed computer systems, that are inspired by the agent paradigm.

One trend is to refine the idea of *object-oriented programming*. The main difference is that agents are *autonomous* in their decisions (whether or not to invoke a particular method, for instance).

Agent-oriented Programming

The central idea in the *agent-oriented programming paradigm* put forward by Shoham (1993) is to program agents in terms of *mentalistic notions* (beliefs, desires, intentions).

Example for a commitment rule:

```
(COMMIT (?a REQUEST ?action)    // message condition
        (B (now (myfriend ?a))) // mental condition
        (?a ?action) )
```

▶ Of course, one may wonder whether it makes sense to ascribe mentalistic notions to a computer system? (Answer: *depends*)

Y. Shoham. *Agent-oriented Programming*. Artif. Intelligence, 60(1):51–92, 1993.

Verification of Agent Systems

The usual issues of *specification and verification* come up in the context of designing both single agents and multiagent systems.

Some of the topics people are working on include:

- Design of new logics to model individual agent behaviour or entire multiagent systems
- Model checking algorithms for BDI logics, or other logics deemed appropriate for modelling intelligent agents
- Conformance checking of agent communication with respect to given interaction protocols

Other Topics in MAS Research

- Argumentation-based interaction in multiagent systems
- Trust and reputation
- Distributed constraint satisfaction
- Multiagent learning
- Mobile agents
- Multiagent systems and robotics
- Embodied, emotional and believable agents
- Specialised research communities interested in specific applications
- ...

Background Reading

The following books are good starting points to find out more about multiagent systems in general:

- M. Wooldridge. *An Introduction to Multiagent Systems*. John Wiley and Sons, 2002.
- G. Weiss (ed.). *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, 1999.

Part II

Scenario

In very general terms, we are going to be interested in the following type of problem:

- Allocate a set of *goods* (or tasks) amongst several *agents*.
- The agents should play an active role in the *allocation procedure*.
- Their actions may be influenced by their individual *preferences*.

Let's first look at some of the possible instantiations of this general problem, for different kinds of applications ...

The Postmen Domain

Our agents are (two) postal workers. They meet in the post office in the morning and discuss the fact that Ann has letters for Carol, Dick and Ed, while Bob has letters for *Frank*, *Gary* and *Hugh* ...

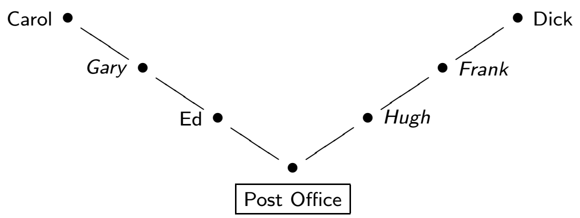

J.S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.

Industrial Procurement

Our agents are suppliers of components for the car manufacturing industry. The "goods" to be allocated are contracts to supply certain numbers of certain components. We want to build 500 cars ...

- Supplier A is selling tyres in packs of 50 for €1000 a pack. If you take more than 20 packs, you get a 17.5% discount.
- Supplier B is selling them in packs of 60, also for €1000 a pack. He can offer no discounts, but anyone buying tyres for over €10000 can benefit from the special offer on engines.
- And so on ...

Industrial Procurement (cont.)

We may also want to take factors other than money into account:

- *Non-price attributes of goods:* quality, time to delivery, ...
- *Constraints over multiple goods:* the colour of the cars does not matter, as long as they do not have all the same colour and as long as there are at least 50 of each colour.
- *Attributes of suppliers:* qualification, certification, past performance, trust and reputation, ...
- *Safety constraints:* don't rely on just a single agent to supply *all* the tyres.

A. Giovannucci et al. *Towards Automated Procurement via Agent-aware Negotiation Support*. Proc. AAMAS-2004.

Earth Observation Satellites

Our agents are representatives of different European countries that have jointly funded a new Earth Observation Satellite (EOS). Now the agents are requesting certain photos to be taken by the EOS, but due to physical constraints not all requests can be honoured ...

Allocations should be both *efficient* and *fair:*

- The satellite should not be underexploited.
- Each agent should get a return on investment that is at least roughly proportional to its financial contribution.

M. Lemaître, G. Verfaillie, and N. Bataille. *Exploiting a Common Property Resource under a Fairness Constraint: A Case Study*. Proc. IJCAI-1999.

Settling Divorce Disputes

Our agents used to be happily married, but have fallen out of love. How should they divide their belongings?

- I value this carpet at €2000. (*quantitative preference*)
- I'd rather have the red than the blue car. (*ordinal preference*)
- No way he's gonna get the piano! (*externalities*)
- I won't be content with less than what she gets! (*envy*)
- I'd rather kill the dog than let him have it! (*pure evil*)

S.J. Brams and A.D. Taylor. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996.

First Approach: A Sealed-Bid Auction

Suppose we have just one item to be allocated.

Procedure: Each agent sends a letter stating how much they would be prepared to pay for the item to an independent "auctioneer". The auctioneer awards the item to the agent bidding the highest price and takes the money.

At least if all agents are honest and send in their *true valuations* of the item, then this procedure has some appealing properties:

- The procedure is *simple*, both computationally and in terms of its communication requirements.
- The agent who likes the item the most will obtain it.
- The auctioneer will make maximum profit.
- No other solution would be better for some of the participants without being worse for any of the others (*Pareto optimality*).

Strategic Considerations

But what if the agents are not honest . . . ?

What would be the best possible strategy for a *rational* (selfish) agent? That is, how much should an agent bid for the item on auction?

▶ The procedure actually isn't simple at all for the agents, and our nice properties cannot be guaranteed.

▶ Try to set up an allocation mechanism giving agents an incentive to bid truthfully: *mechanism design* ($\leadsto$ *game theory*).

Auctioning Multiple Items

Strategic issues aside, our simple auction mechanism is not that bad for allocating a single item. But what if there are *several goods*?

Suppose the auctioneer is first selling a TV and then a DVD player. How should our agents bid?

- Ann wants to watch the news and is not interested in DVDs.
- Bob already owns a TV and is only interested in the DVD player.
- Chloë has an enormous collection of classic movies on DVD, a pretty low opinion of today's television programming, and no DVD player or TV.

Complements and Substitutes

The value an agent assigns to a *bundle* of goods may relate to the value it assigns to the individual goods in a variety of ways ...

- *Complements:* The value assigned to a set is *greater* than the sum of the values assigns to its elements.

  A standard example for complements would be a pair of shoes (a left shoe and a right shoe).

- *Substitutes:* The value assigned to a set is *lower* than the sum of the values assigned to its elements.

  A standard example for substitutes would be a ticket to the theatre and another one to a football match for the same night.

In either case a simple auction mechanism that allocates one item at a time is problematic, even if we were to make the (unrealistic) assumption that agents will bid truthfully ...

Combinatorial Auctions

In a *combinatorial auction*, the auctioneer puts several goods on sale and the other agents submit bids for entire bundles of goods.

Given a set of bids, the *winner determination problem* (WDP) is the problem of deciding which of the bids to accept.

- The solution must be *feasible* (no good may be allocated to more than one agent).
- Ideally, it should also be *optimal* (in the sense of maximising revenue for the auctioneer).

Clearly finding a solution to the WDP can be tricky (just how "tricky" we'll see later on in the course).

So besides the game-theoretical problem of stopping bidders from strategising, in combinatorial auctions we also face a challenging *algorithmic* problem.

Communicating Bids

Suppose we are running a combinatorial auction with n goods. So there are $2^n - 1$ bundles that agents may want to bid for.

For interesting values of n, it is not possible to communicate your valuations to the auctioneer by simply stating your price for each and every bundle. $\leadsto$ How do we best communicate/represent preferences in combinatorial domains?

For combinatorial auctions, this is the job of the *bidding language*.

Example: Bid for a small number of concrete bundles with the implicit understanding that you will honour any combination of bids that is feasible and pay the sum of the associated prices.

In general, *preference representation* is a central issue in multiagent resource allocation ...

Preference Representation Languages

- *Cognitive relevance:* How close is a given language to the way in which humans would express their preferences?
- *Elicitation:* How difficult is it to elicit the preferences of an agent so as to represent them in the chosen language?
- *Expressive power:* Can the chosen language encode all the preference structures we are interested in?
- *Succinctness:* Is the representation of (typical) preference structures succinct? Is one language more succinct than the other?
- *Complexity:* What is the computational complexity of related decision problems, such as comparing two alternatives?

Centralised vs. Distributed Approaches

So far we have concentrated on auctions as mechanisms for solving resource allocation problems. But what if we cannot find someone who could act as the auctioneer?

- The associated tasks may be too hard *computationally* for the agent supposed to assume the role of auctioneer.
- There may be no agent that enjoys the *trust* of the others.
- The system infrastructure may be truly *distributed* and it may simply be unnatural to model the problem in a centralised manner.
- The goods may be owned by different agents to begin with (that is, we may have to take an *initial allocation* into account).
- Agents and goods may enter or leave the system *dynamically*.

Auctions are *centralised* mechanisms. All of the above are good reasons to consider *distributed* negotiation schemes as well . . .

Distributed Negotiation Schemes

In truly distributed approaches to resource allocation, allocations emerge as the result of a sequence of local negotiation steps.

While this is closer to the true spirit of multiagent systems, such systems are also more difficult to design and understand than auction-based mechanisms. Issues include:

- What types of deals do we allow for (*bilateral* vs. *multilateral* negotiation)?
- How do we design appropriate communication *protocols*? Note that this problem is relatively easy for auctions.
- To what degree is it possible to *predict* (or even control) the outcome of distributed negotiation processes? Put differently: What is the relationship between the negotiation strategies agents use *locally* and the allocations emerging *globally*?

What is a "good" allocation?

In the case of an auction we can measure the quality of an allocation from the viewpoint of the auctioneer: the more *revenue* the better.

For distributed negotiation schemes, or if the auction is just a means to an end, the quality of an allocation should *somehow* depend on the *individual preferences* of the agents in the system. But how?

Multiagent systems may be thought of as *"societies of agents"* . . .

The relationship between the preferences of individuals and the well-being of society as a whole (*social welfare*) has also been studied in *social choice theory*.

Efficiency and Fairness

When assessing the quality of an allocation (or any other agreement) we can distinguish (at least) two types of indicators of social welfare.

Aspects of *efficiency* (*not* in the computational sense) include:

- The chosen agreement should be such that there is no alternative agreement that would be better for some and not worse for any of the other agents (*Pareto optimality*).
- If preferences are quantitative, the sum of all payoffs should be as high as possible (*utilitarianism*).

Aspects of *fairness* include:

- The agent that is going to be worst off should be as well off as possible (*egalitarianism*).
- No agent should prefer to take the bundle allocated to one of its peers rather than keeping their own (*envy-freeness*).

Complexity Issues

It should be clear by now that it is certainly not easy to set up a good mechanism for resource allocation in a multiagent system.

But what about the precise *computational complexity* of finding an optimal allocation?

► While the problems we are interested in originate in economics, computer science (amongst others) can provide us with very useful tools for the formal analysis of these problems.

A Different Perspective on the Course

This course will address certain issues in Multiagent Systems research.

But it may also be characterised as a course on topics at the interface of Artificial Intelligence and Computer Science on the one hand, and the Socio-economic Sciences on the other.

People keep inventing new names for this emerging field:

- "Computational Game Theory",
- "Computational Social Choice",
- "Social Software",
- ...

Related Courses

- Game Theory for Information Sciences (last semester)
  Peter van Emde Boas
- Multiagent Systems and Distributed AI (MSc AI)
  Nikos Vlassis
- Caput Logic, Language and Information (focus on Social Software)
  Eric Pacuit
- A new course in the next academic year
  Krzysztof Apt

What next?

Over the next couple of weeks or so we'll be going through some relevant background material:

- Social Choice and Welfare
- Game Theory

This will provide us with some of the basic formal tools required to analyse the problems presented today in more detail.