



EH2750 Computer Applications in Power Systems, Advanced Course.

Lecture 8

Professor Lars Nordström, Ph.D.
Dept of Industrial Information & Control systems, KTH
larsn@ics.kth.se



Acknowledgement

- These slides are based largely on a set of slides provided by:

*Professor Rosenschein of the Hebrew University
Jerusalem, Israel*

and

Dr. Georg Groh, TU-München, Germany.

- Available at the Student companion site of the Introduction to Multi Agent Systems book
-




Outline of the Lecture

- Repeating where we are right now
 - Intelligent Agents of various types
 - How to make agents think and plan
 - Multi-agent interactions
 - Some concepts for cooperation
 - Allocating Scarce Resources - Auctions
-



What is an Intelligent Agent?

- The main point about agents is they are *autonomous*: capable of acting independently, exhibiting control over their internal state
- Thus: *an intelligent agent is a computer system capable of flexible autonomous action in some environment in order to meet its design objectives*






The discussion so far

- Chapter 2 describes the idea of agents that perform tasks in an environment and sets some definitions
- Chapters 3, 4, & 5 describe three different approaches to describing and developing the apparent Intelligence in the agents.
 - Chapter 3 – Deductive Reasoning Agents
 - Chapter 4 – Practical Reasoning Agents
 - Chapter 5 - Reactive (and Hybrid Agents)
- In the Excerpt from the AI book used in Lecture #4 we took a look at planning and searching
- Today we start looking at the Multi in Multi-agent systems



Practical Reasoning

- Human practical reasoning consists of two activities:
 - *deliberation*
deciding *what* state of affairs we want to achieve
 - *means-ends reasoning*
deciding *how* to achieve these states of affairs
- The outputs of deliberation are *intentions*





Practical Reasoning Agent

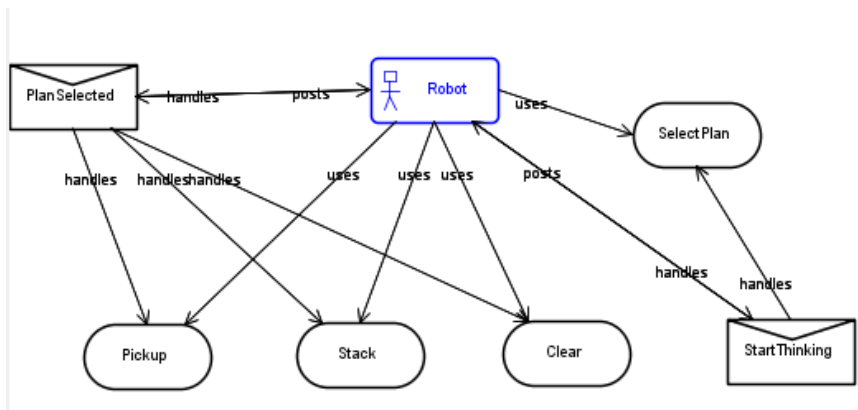
```

function SIMPLE-PROBLEM-SOLVING-AGENT(percept) returns an action
inputs: percept, a percept
static: seq, an action sequence, initially empty
         state, some description of the current world state
         goal, a goal, initially null
         problem, a problem formulation

state ← UPDATE-STATE(state, percept)
if seq is empty then do
    goal ← FORMULATE-GOAL(state)
    problem ← FORMULATE-PROBLEM(state, goal)
    seq ← SEARCH(problem)
action ← FIRST(seq)
seq ← REST(seq)
return action
    
```



How this can look in JACK



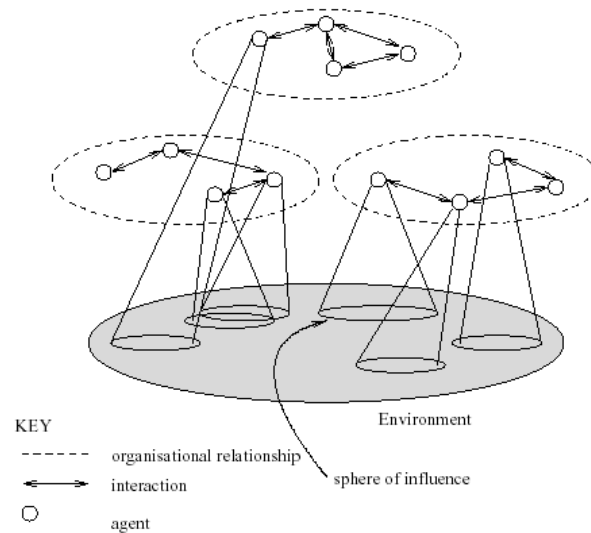


Outline of the Lecture

- Repeating where we are right now
 - Intelligent Agents of various types
 - How to make agents think and plan
- Multi-agent interactions
 - Some concepts for cooperation
- Allocating Scarce Resources - Auctions



Multi-agent Systems





Multi-agent Systems

Contains a number of agents...

- ...which interact through communication...
 - ...are able to act in an environment...
 - ...have different “spheres of influence” (which may coincide)...
 - ...will be linked by other (organizational) relationships
-



Working Together

- Why and how do agents work together?
 - Important to make a distinction between:
 - *benevolent agents*
 - *self-interested agents*
-



Benevolent Agents

- 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)
 - *Benevolence simplifies the system design task enormously!*
-



Self-Interested Agents


- If agents represent individuals or organizations, (the more general case), then we cannot make the benevolence assumption
 - Agents will be assumed to act to further their own interests, possibly at expense of others
 - Potential for *conflict*
 - May complicate the design task enormously
-



Benevolent Agents

Task Sharing and Result Sharing


- Two main modes of cooperative problem solving:
 - *task sharing*: components of a task are distributed to component agents
 - *result sharing*: information (partial results, etc.) is distributed



Benevolent Agents

Cooperative Distributed Problem Solving

- CDPS is concerned with investigation of:
 - Problem subdivision
 - Sub-Problem distribution
 - Result synthesis
 - Optimization of problem solver coherence
 - Optimization of problem solver coordination





Benevolent Agents Coherence

Coherence: Refers to “how well the MAS behaves as a unit along some dimension of evaluation”. Coherence may be measured in terms of

- Solution quality
- resource usage
- conceptual clarity of operation
- performance degradation if unexpected failure occurs



Benevolent Agents Coordination

• **Coordination:** “The degree...to which [the agents] can avoid ‘extraneous’ activity [such as] ...synchronizing and aligning their activities”

→ Poor coordination if

- Agents clobber each other’s sub-goals
- Lots of communication (no mutual predictability (e.g. by expressive models of each other))
- Destructive interference if conflict



Self-Interested Agents



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*:

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

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

- Utility functions lead to *preference orderings* over outcomes:

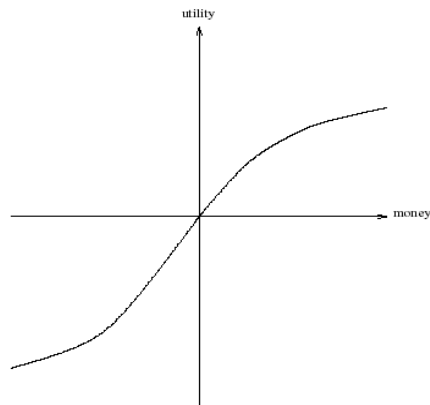
$$\omega \succeq_i \omega' \text{ means } u_i(\omega) \geq u_i(\omega')$$

$$\omega \succ_i \omega' \text{ means } u_i(\omega) > u_i(\omega')$$



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:

$$u_i(\omega_1) = 1 \quad u_i(\omega_2) = 1 \quad u_i(\omega_3) = 4 \quad u_i(\omega_4) = 4$$

$$u_j(\omega_1) = 1 \quad u_j(\omega_2) = 4 \quad u_j(\omega_3) = 1 \quad u_j(\omega_4) = 4$$

- With a bit of abuse of notation:

$$u_i(D, D) = 1 \quad u_i(D, C) = 1 \quad u_i(C, D) = 4 \quad u_i(C, C) = 4$$

$$u_j(D, D) = 1 \quad u_j(D, C) = 4 \quad u_j(C, D) = 1 \quad u_j(C, C) = 4$$

- Then agent i 's preferences are:

$$C, C \succeq_i C, D \succ_i D, C \succeq_i D, D$$

- “C” is the *rational choice* for i .
(Because 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	4
	coop	1	4

- Agent i is the *column player*
- Agent j is the *row player*



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



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*
-



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:

$$u_i(\omega) + u_j(\omega) = 0 \quad \text{for all } \omega \text{ in } \Omega$$
 - 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
-



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



The Prisoner's Dilemma

- Payoff matrix for prisoner's dilemma:

		<i>i</i>	
		defect	coop
<i>j</i>	defect	2	1
	coop	4	3
		1	3

- Top left: If both defect, then both get punishment for mutual defection
- Top right: If *i* cooperates and *j* defects, *i* gets sucker's payoff of 1, while *j* gets 4
- Bottom left: If *j* cooperates and *i* defects, *j* gets sucker's payoff of 1, while *i* gets 4
- Bottom right: Reward for mutual cooperation



The Prisoner's Dilemma

- 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;
 - The prisoner's dilemma is *present everywhere*.
 - Can we recover cooperation?
 - Well, yes we can introduce auctions, negotiations and argumentation. More on this next lecture!
-



Outline of the Lecture

- Repeating where we are right now
 - Intelligent Agents of various types
 - How to make agents think and plan
 - Multi-agent interactions
 - Some concepts for cooperation
 - Allocating Scarce Resources - Auctions
-



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, . . .)
 - Network capacity
 - Amount of energy
 -
-



Reaching Agreements

- The extreme case of a Multiagent encounter is the zero-sum. (*Profit only at expense of others*)
- Normal case is the "Win-win-situation" where mutually beneficial agreement is possible
- Reaching Agreements is fundamental for social intelligence and society building in general
- Reaching Agreements is a result of Negotiation, Auctions and/or Argumentation
- How do you construct algorithms these types of interactions



Algorithm Design Criteria

- We want to have a design of the algorithm that has certain properties:
 1. Guaranteed success - Agreement is certain
 2. Maximizing social welfare - Agreement maximizes sum of utilities of all participating agents
 3. Computationally efficient



Algorithm Design Criteria, continued

4. Pareto efficiency: iff there exists no other agreement which increases utility of at least one agent while not decreasing the utility of the other agents
5. Individual Rationality: Following protocol is in best interest of all agents (no incentive to cheat, deviate from protocol etc.)
6. Stability: Protocol gives agents incentive to behave in a certain way. (→ e.g. by establishing Nash-Eq.)
7. Simplicity: Protocol makes for the agent appropriate strategy „obvious“. (Agent can tractably determine optimal strategy)
8. Distribution: no single point of failure; minimize communication



What are auctions?

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



Auctions

- Auctions are simple → easy to implement
- Auction = consists of (Auctioneer, Bidders, Good);
Goal of the Auctioneer is to maximize price for good;
Goal of the Bidders is to minimize price for good;
Each bidder has personal price maximum
- Auctioneer: Tries to reach goal by choosing appropriate auction mechanism
- Bidders: Try to reach goal by choosing appropriate strategy
- Auction algorithms differ by:
 - Winner determination,
 - Secrecy of bids,
 - Auction procedure



Single vs. Multi-dimensional auctions

- 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."



Value of the goods

Good has a public (common) value: Good has the same value for all bidders. (E.g.: One-Dollar-Bill)

Good has private value: Good has different value for each agent. (E.g.:)

Good has correlated value: Value of good depends on own private value and private value for other agents. (E.g.: Buy sth. with intention to sell it later)



Winner Determination

- First price: Highest bid wins, Winner pays his bid
- Second price: Highest bid wins. Winner pays second-highest bid
- General case, highest bid wins, pays $n-k$ bid.



Secrecy of the Bids

- Open cry: All agent's know all agent's bids.
- Sealed bid: No agent knows other agent's bids



Auction Procedure

- One shot: Only one bidding round
- Ascending
 - Auctioneer begins at minimum price, bidders increase bids
 - Also known as English Auction
- Descending
 - Auctioneer begins at price over value of good and lowers the price at each round
 - Also known as Dutch auction



English Auctions

- Most common form (human world)
- **Open cry, first price, ascending.**
- Dominant strategy: Bid slightly more than current bit, withdraw if bid reaches personal valuation of good
- If uncertainty of (private or public) value of good exists:
 - “Should you be happy that you won the good?”
 - “Why did the other bidders not bid more?”
- Possibly: Although winning bid is below personal valuation of winner, the “true value” may be less than bid → “Winner’s curse”. (E.g. bidding for gold-mine)



Dutch Auctions


- **Open cry, first-price, descending.**
- No dominant strategy. Winner’s curse also possible.





Sealed-bid First-price One-shot

- "Pseudo"-dominant strategy: (Assuming that others bid their true valuation): Bid less than true valuation.




Vickrey Auctions

- **Sealed-bid, second-price, one shot**
- Dominant strategy: (Assuming that others bid their true valuation): Bid true valuation.

Why? If you bid less than your true valuation \rightarrow you only decrease your chances, but you will not influence the price you have to pay.

(As in all auctions: If you bid more than your true valuation \rightarrow you risk winner's curse)





Combinatorial Auctions

- Until now, we have considered auctions for one indivisible good.
 - Now consider auctions for goods that are divisible
 - Like for instance
 - Amount of available energy next hour, or
 - Power that can flow along a line
-



Formulation

$$\mathcal{Z} = \{z_1, \dots, z_m\}$$

is a set of items to be auctioned, we have the usual set of agents

$$Ag = \{1, \dots, n\},$$

and we capture preferences of agent i with the valuation function:

$$v_i : 2^{\mathcal{Z}} \mapsto \mathbb{R}$$

meaning that for every possible bundle of goods

$$Z \subseteq \mathcal{Z}, v_i(Z)$$

says how much Z is worth to Agent i .



Some additional facts & assumptions

If

$$v_i(\emptyset) = 0$$

then we say that the valuation function for i is normalised. Meaning that if an agent is allocated nothing it is worth nothing.

Similarly, we have that

$$Z_1 \subseteq Z_2 \text{ implies } v_i(Z_1) \leq v_i(Z_2)$$



Winner (or allocation) determination

- An allocation is a list of sets (allocations) Z_1, \dots, Z_n for each agent Ag_i so that

$$Z_i \subseteq \mathcal{Z}$$

- And for all i, j in Ag such that $i \neq j$ we have $Z_i \cap Z_j = \emptyset$
- Valid for discrete sets of goods

- A general continuous case is similarly constrained by the sum of $Z_1, \dots, Z_n \leq Z$



How to do the allocation then?

- A reasonable assumption is to allocate in a way that maximises the social welfare, i.e. Maximizing the total value achieved, the sum of all utilities.

$$sw(Z_1, \dots, Z_n, v_1, \dots, v_n) = \sum_{i=1}^n v_i(Z_i)$$



Combinatorial Auction setup

- Given this, we can define a *combinatorial auction*.
- Given a set of goods \mathcal{Z} and a collection of valuation functions v_1, \dots, v_n , one for each agent $i \in Ag$, the goal is to find an allocation

$$Z_1^*, \dots, Z_n^*$$

that maximizes sw , in other words

$$Z_1^*, \dots, Z_n^* = \arg \max_{(Z_1, \dots, Z_n) \in alloc(\mathcal{Z}, Ag)} sw(Z_1, \dots, Z_n, v_1, \dots, v_n)$$

- Figuring this out is *winner determination*.



Determining the allocation

- How do we do this?
 - Well, we could get every agent i to declare their valuation \hat{v}_i
 - The hat denotes that this is what the agent *says*, not what it necessarily is.
 - The agent may lie!
 - Then we just look at all the possible allocations and figure out what the best one is.
-



Computational efficiency?

- One problem here is *representation*, valuations are exponential:

$$v_i : 2^Z \mapsto \mathbb{R}$$
 - A naive representation is impractical.
 - In a bandwidth auction with 1122 licenses we would have to specify 2^{1122} values for each bidder.
 - Searching through them is computationally intractable.
-



So, how do we do it then?

- Searching through all combinations is a basic problem but intractable due to computation resources needed.
 - However, this is the worst case result, so it may be possible to
 - We can try to develop approaches that are optimal and run well in many cases.
 - Can also forget optimality and either:
 - use heuristics; or
 - look for approximation algorithms.
 - Common approach: code the problem as an integer linear program and use a standard solver – often works in practice.
 - In practice a constraint satisfaction problem, that can be solved with different search mechanisms
-



Outline of the Lecture

- Repeating where we are right now
 - Intelligent Agents of various types
 - How to make agents think and plan
 - Multi-agent interactions
 - Some concepts for cooperation
 - Allocating Scarce Resources - Auctions
-