Human-like artificial creatures

9. Intelligent autonomous agents

Cyril Brom
Faculty of Mathematics and Physics
Charles University in Prague
brom@ksvi.mff.cuni.cz

Outline

1. Definition of agency
   - Agent, intelligent agent, autonomy
   - IVA, animat, virtual human
2. Abstract architectures
   - subsumption
   - BDI
   - layered
3. Multi-agents system
   - FIPA
   - speech acts

Autonomous agent

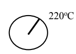

- An **agent** is a computer system that is <u>situated</u> in some environment, and that is capable of <u>autonomous</u> actions in this environment in order to meet its design objectives. [Wooldridge, 1995]
  - a thermostat
  - a software daemon
- An **autonomous agent** is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future. [Franklin, 1997]
  - a thermostat
  - a software daemon
  - an animal

Autonomy

- Knowledge about the environment comes [to the agent] only through the use of the sensor and motor aparatuses, each acting as a kind of communication channel. [Wilson, 1991]
  - Plausibility of sensors and effectors?
- Determined by its own experience.
- A truly autonomous intelligent agent should be able to operate successfully in a wide variety of environments, given sufficient time to adapt. [Russell a Norwig, 1995, p. 35]

Autonomous <u>intelligent</u> agent

- When do we consider an agent to be intelligent?
  - **reactivity** (perception and ability to respond in a timely fashion in order to satisfy design objectives)
  - **proactivness** (goal-directed behaviour by <u>taking the initiative</u> in order to satisfy design objectives)
  - **social ability** (capability of interacting with other agents in order to satisfy design objectives)
- Examples
  - thermostat? unix daemon?
  - web broker, bot,…

[Wooldridge, 1995]

Agents are not objects!

- Passive vs. active
- Autonomy
- Black-box, we cannot prove what the agent will do
- "Objects do it for free, agents do it for money"
- Implementation of agents
  - a program ("heavy" agent)
  - a thread
  - a shared thread ("light" agent)
- Of course, an agent is typically implemented by means of OOP
- The agent oriented programming is a concept, and architecture

  (consider client-server metaphor)

Classification of agents

- autonomous agents
  - biological agents
  - robotic agents
  - computational agents
    - software agents
      - task-specific agents
      - entertainment agents
      - viruses
    - artificial life agents [Franklin, 1997]
- note: consider Minsky's The Society of Mind [Minsky, 1986]

Intelligent virtual agents

- **Intelligent virtual agent** (IVA) is autonomous, graphically embodied agent in an interactive, 2D or 3D virtual environment. It is able to interact intelligently with the environment, other IVAs, and especially with human users.
  - [IVA conferece, 1999 – 2005]

Animat

- originally, Wilson (1991)
  - From animals to animats: Simulation of Adaptive Behaviour
- „**Animat** is a synthetic creature living within a virtual environment that is embodied that means it interacts with the world using only its body, which is subject to the constraints of its environment." [Champandard, 2003, p. 19, 21]
- Originally, biological plausibility was important [Wilson, 1991]:
  - ...the basic **hypothesis** of the **animat approach** is that by simulating and understanding complete animal-like system at a simple level, we can build up gradually to the human.
  - We hope **to reach human intelligence »from bellow«**.
  - ...it is vital to ... maximize physicality in the sensory signals, so as to avoid predefined symbolic inputs.

Our approach – virtual being

- Virtual being (human) is a software, **behavioural model** of an animal (human)
  - it is essentially **embodied**, that means its body is subject to constraints of the environment
  - its environment is a model of the **natural environment** of the creature at an appropriate level of abstraction
    - the environment is typically **dynamic**, **nondeterministic**, and **partially observable**
    - the environment as well as the being's body is typically portrayed **graphically**
  - it carries out sequences of **actions**
    - an action lasts from tens of miliseconds to seconds according to the level of abstraction
    - an action is typically thought as a logical unit
    - an action is a behavioural model of a process, neither a model of a physiological process, nor a model of a mental process.
  - is **complex**
    - must arbitrate between potentially conflicting goals or tasks
    - its body is a model of a whole body (and not a hand for example)

Artificial environments

- accessible/inaccessible
  - an agent can not obtain accurate up-to-date information about the whole environment
- deterministic/non-deterministic
  - the outcome of some actions is not uniquely defined
- static/dynamic
  - the environment changes in ways beyond the agent's control
- discrete/continuous in time/space:
  - finite number of discrete states is garanted
- real-time/step-based
  - the agent has theoretically infinite time to make a decision

[Russell and Norwig, 1995]

Outline

1. Definition of agency
   - Agent, intelligent agent, autonomy
   - IVA, animat, virtual human
2. **Abstract architectures**
   - subsumption
   - BDI
   - layered
3. Multi-agents system
   - FIPA
   - speech acts

Abstract architectures for intelligent agents (1)

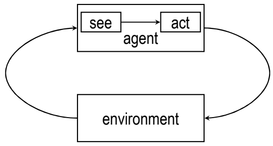

- Environment may be in any of a finite state E: **E** = {e_1, e_2, …}
- Agent's actions transform the state of the environment. Let **A** = {a1, a2, …} be a finite set of actions (environment is non-deterministic)
- Let $\mathbf{R}^E$ be the subset of all possible finite sequences (over E and A) that start in the e_0 and end in an e_n: $e_0 \xrightarrow{a_1} e_1 \xrightarrow{a_2} e_2 \xrightarrow{a_3} \dots \xrightarrow{a_{n-1}} e_n$
- **Deterministic agent** is a function that maps the history of the system to actions: ***Ag***: $R^E \rightarrow A$

Abstract architectures for intelligent agents (2)

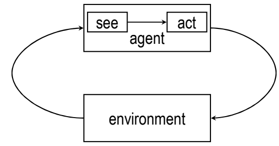

- **Pure reactive** agent can be represented by a function: ***Ag***: E → A (e.g. a thermostat)
- Sensing: Let a percept be a perceptual input. Let **Per** be a non-empty set of percepts:
  - **see**: E → Per
  - **act**: Per* → Ac (Per → Ac, for a reactive agent, respectively)
- If two different environmental states are mapped to the same percept, the agent will not be able to distinguish between them
- Pure reactive agents reject of symbolic representation and of decision making based on syntactic manipulation of such representations

Subsumption architecture

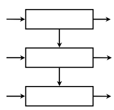

- Represented by sets of behaviours (c, a)
- *c* is a set of percept (condition), *a* is an action
- We say, that a behaviour fire iff see(s)∈c
- There is a total ordering on the set of behaviours. If b_1 is "lower" in the hierarchy then b_2, we say that b_1 inhibits b_2.
- The action of the lowest firing behaviour is chosen to execute

[Brooks, 1986]
[Wooldridge, 1995]

Subsumption architecture
implementation

- Originally implemented in a robot, every layer was realized as a hardwired finite-state-machine
- It can be implemented by means of if-then rules:

```
1. if see_obstacle then change_direction
2. if basketful_of_m. and picking then stop_picking
3. if see_m. and picking then pick_up_The_m.
4. if midday and picking then stop_picking
5. if home then END
6. if picking then move_random
7. if not_picking then move_home
```

But is this example pure subsumption, or not?

Abstract architectures for intelligent agents (3)

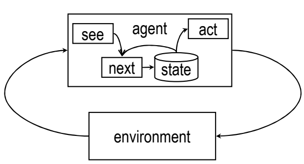

- Agents with **state**: Let ***I*** be a set of all possible internal states of the agent
  - **see**: E → Per
  - **next**: I x Per → I
  - **act**: I → Ac
- How to make decisions? Utility value represents how "good" a state is:
  - **u**: E → Real (utility of states)
  - **u**: R^E → Real (utility of runs)

Practical reasoning - BDI

- A model of decision making. Practical reasoning is reasoning directed towards actions
  - what to do (**deliberative** reasoning)
  - how to do it (**means-ends** reasoning)
- The state of a BDI agent in any given moment is represented by (Bel, Des, Int).
- **Belief revision function** (*brf*) takes percepts and the agent's current beliefs and on their basis, determines a new set of beliefs:
  - *brf*: P(Bel) x Per → Bel
- **Option generation function** (*opt*) generates new desires according to the current agent's intentions and the current beliefs
  - *opt*: P(Bel) x P(Int) → Des
- **Filter function** (*fil*) sets "the best option" for the agent to commit to; it determines the agent's intentions
  - *fil*: P(Bel) x P(Int) x P(Des) → P(Int)
- **Action selection function** (*exec*) determines, which action to perform on the basis of current intentions
  - *exec*: P(Int) → Ac

[Bratman, 1987]
[Wooldridge, 1995]

Abstract architectures for intelligent agents (4)

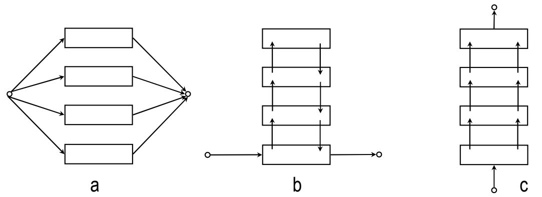

- Hybrid approach: how to combine reactivity with planning capabilities?
- Layered architectures
  - horizontally (a)
  - vertically (b, c)
- Three-layered architectures
  - reactive layer
  - deliberative layer (planning)
  - reflexive meta-level reasoning (planning about planning)

Outline

1. Definition of agency
   - Agent, intelligent agent, autonomy
   - IVA, animat, virtual human
2. Abstract architectures
   - subsumption
   - BDI
   - layered
3. **Multi-agents system**
   - FIPA
   - speech acts

A multi-agent system I

- Agents can meet each other
- An agent can be created or destroyed
- An agent can migrate between systems
- An agent might want to let the other agents know what (s)he can do for him
- An agent might want to know what the other agents can do for him(her)

- Essentially peer-to-peer

A multi-agent system II

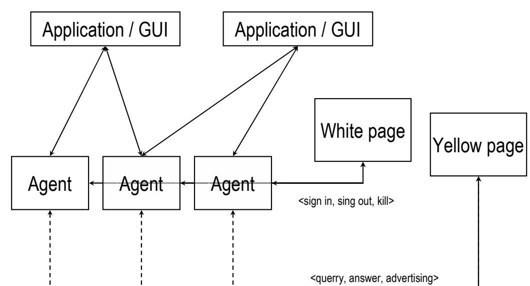

A multi-agent system III

- White page
  - an agent administrator
  - holds the real information about the agent (i.e. id)
  - authority
    - identification, create, kill,...
- Yellow page
  - a special-purpose agent, advertising service
  - querying, subscribing
  - a yellow-page federation

A multi-agent system IV

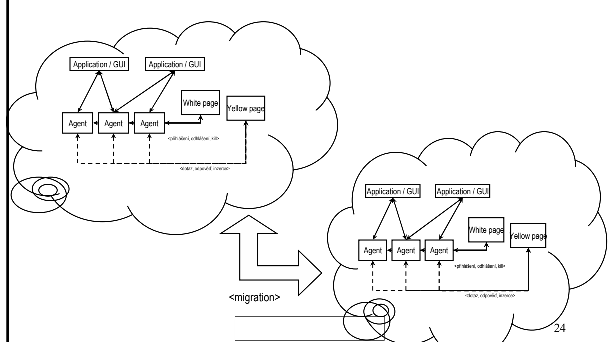

FIPA agents specification

- The Foundation for Intelligent Physical Agents
- A specification for multi-agent systems and agent's communication
  - the structure of a multi-agent system
  - communication language
  - interaction protocols
- More specifications exist (ie. KQML)

A platform

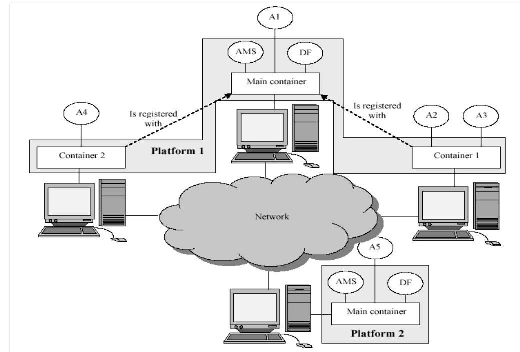

[FIPA specification] (c)

FIPA
communication I

- FIPA0037 – FIPA communicative Act Library Specification
- layers
  - physical layer (wiring)
  - transport layer (TCP/IP, ...)
  - communication layer (CORBA, JADE, ...)
  - ACL - agent communication language (**FIPA-ACL**, KQML)
  - message language (**FIPA-SL**, Prolog, XML, ...)

FIPA
communication II – FIPA ACL

- Based on speech act theory
  - **speech acts** – a class of natural language utterances that have the characteristic of actions – they change the state of the world in a way analogous to physical actions
    - "I now pronounce you man and wife."
    - a note: this approach is also used in the domain of virtual humans
  - **performative verbs** – they correspond to different types of speech acts
    - inform, promise, request,..
- History
  - speech acts: Austin, 1962, *How to do things with words*
  - detailing, preconditions, classification: Searle, 1969, *Speech Acts*
  - formalisation: "...modelling speech acts in a planning system as operators defined... in terms of speakers' and hearers' beliefs and goals": Cohen, Perrault, 1979

FIPA
communication III – speech acts in FIPA ACL

- About 20 communicative acts (performatives)
- The semantics of the FIPA ACL map each ACL message to a formula of SL (semantic language)
  - SL defines constrains on the sender and receiver
  - based on modal logic and possible-worlds-semantics: $\Box\varphi$, $\Diamond\varphi$
  - ontologies for agents communication

FIPA
communication IV

```
(cfp
  :sender (agent-identifier :name j)
  :receiver (set (agent-identifier :name i))
  :content
    "((action (agent-identifier :name i)
       (sell plum 50))
      (any ?x (and (= (price plum) ?x) (< ?x 10))))"
  :ontology fruit-market
  :language fipa-sl)
```

Perform the following: send me a message whether or not you sell me 50 kg of plums and for how much is a kilo – but I don't accept more than 10$ for a kilo.

- **What kind of information is specified by the ontology?**
- **Can you imagine a different ontology?**

FIPA

communication V

- Contract Net – FIPA0029
- Eg. find the cheapest book

How many messages are sent totally if *m* is the number of addresseds, from which *n* addresseds answer?

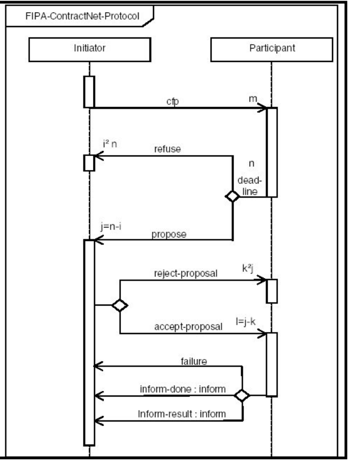

Applications

- Mobile networks
  - JADE framework
- E-commerce
- Simulations, planning
- Semantic web
  - web-services?, grids?

End.

References

1. S. Franklin, A. Graesser: Is it an agent, or just a program? In: *Intelligent Agents III*, LNAI Volume 1193, Springer, Berlin (1997)
2. S. J. Russell and P. Norwig: *Artificial Intelligence: a Modern Approach*. Prentice-Hall, Englewood Cliffs, NJ. (1995, 2003)
3. M. Wooldridge: *An Introduction to MultiAgent Systems*. John Wiley & Sons (1995)
4. Minsky, M.: *The Society of Mind*. Simon and Schuster Inc., (1985)
5. Huber, M. J.: JAM: A BDI-theoretic mobile agent architecture. In: *Proceedings of the 3rd International Conference on Autonomous Agents* (Agents'99). Seatle (1999) 236-243
6. FIPA: www.fipa.org
7. Bratman, M. E.: *Intention, Plans and Practical Reason*. Harvard University Press, Cambridge, Ma (1987)
8. S. W. Wilson. The animat path to AI. In J.-A. Meyer and S. Wilson, editors, From Animals to Animats, pages 15–21. MIT Press, Cambridge, MA, 1991.