Introduction to Goals in Multi-Agent Systems

M. Birna van Riemsdijk, TU Delft, The Netherlands

4/24/11

Overview Goals Tutorial

- Monday
  - introduction to goals
  - representation of goals, goal types
- Tuesday
  - dynamics of goals
  - modularity, interaction, organization
- Thursday
  - practical session with GOAL agent programming language

Combination of theory & practice

1.

My Background

PhD 2002-2006

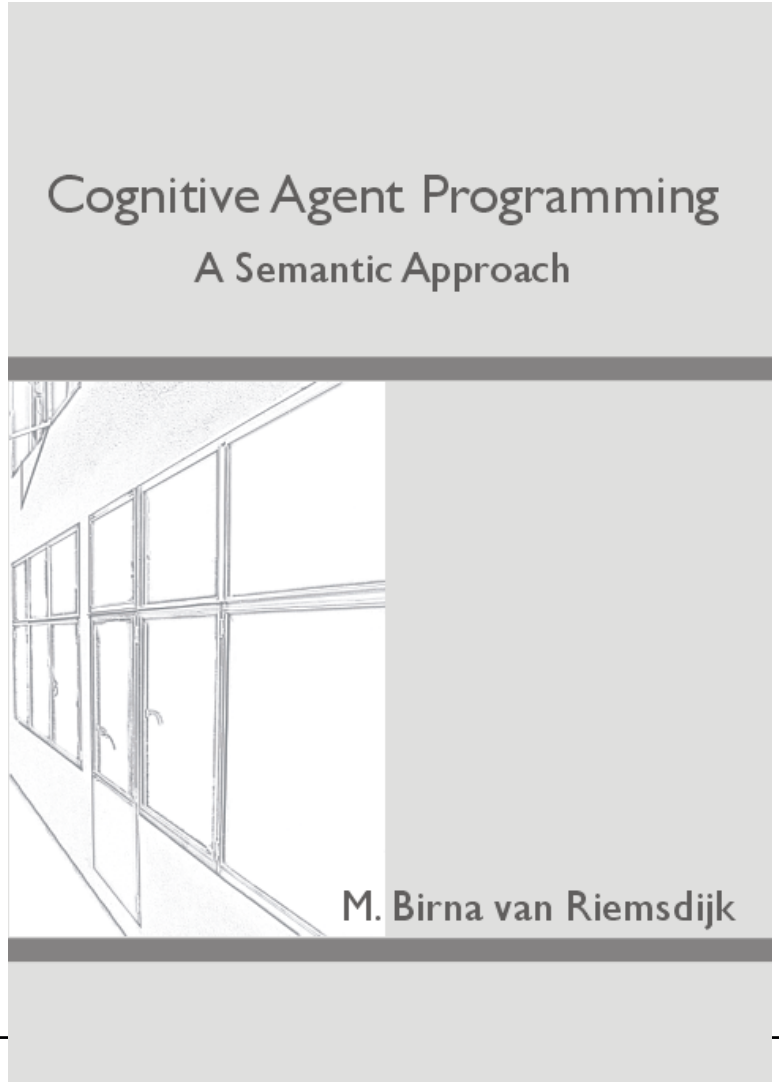

Utrecht University
Intelligent Systems

John-Jules Meyer
Frank de Boer
Mehdi Dastani

Postdoc 2006-2008

Ludwig Maximilians Universitaet Muenchen
Programming & Software Engineering

Martin Wirsing

http://massengale.typepad.com/venustas/2004/09/munich_new_york.html

Delft

Assistant Professor TU Delft 2008-

Agents@MMI: theory & practice

- GOAL agent programming language
  - Formal semantics, verification logic, model checking
  - Empirical software engineering for agent programming
  - Organization-aware agents & shared mental models
- Pocket negotiator

GOAL Agent Programming Language

- http://mmi.tudelft.nl/trac/goal/
- first proposed by Koen Hindriks et al. (ATAL 2000)
- used in teaching since 2007
- comes with IDE

K. V. Hindriks, F. S. de Boer, W. van der Hoek, and J.-J. Ch. Meyer. Agent programming with declarative goals. In IntelligentAgents VI - Proceedings of the 7th International Workshop on Agent Theories, Architectures, and Languages (ATAL'2000),Lecture Notes in AI. Springer, Berlin, 2001.

K. V. Hindriks. Programming rational agents in GOAL. In R. H. Bordini, M. Dastani, J. Dix, and A. El Fallah Seghrouchni,editors, Multi-Agent Programming: Languages, Tools and Applications. Springer, Berlin, 2009.

My DALT History

- Attended 1st DALT: 2003
- Dynamics of declarative goals in agent programming (DALT'04)
- Plan Generation and Plan Execution in Agent Programming (DALT'06)
- Satisfying Maintenance Goals (DALT'07)
- Using Temporal Logic to Integrate Goals and Qualitative Preferences into Agent Programming (DALT'08)
- Co-chair: 2007, 2008, 2009
- DALT Steering committee: 2009-...

2.

Introduction to Goals

What is a Goal?

http://www.thestar.com/sports/soccer/worldcup/article/832712--flashy-win-puts-netherlands-in-world-cup-final

Well...

dictionary.reference.com

goal:
-noun

1. result or achievement toward which effort is directed; aim; end.
2. the terminal point in a race
3. ....

Back to where it all started...

http://www.paintinghere.com/painting/The_Tree_of_Life_1944.html Gustav Klimt, Tree of Life

3.

Philosophy

The Intentional Stance (1987)

Daniell C. Dennett, Philosopher

Predicting & explaining the behavior of an object

- Physical stance: physics & chemistry
- Design stance: biology & engineering
- Intentional stance: software & minds

http://www.crazymalc.co.nz/2007/Dec/24Dec/24Dec.htm

The Intentional Stance (1987)

Daniell C. Dennett, Philosopher

Here is how it works: first you decide to treat the object whose behavior is to be predicted as a rational agent; then you figure out what beliefs that agent ought to have, given its place in the world and its purpose. Then you figure out what desires it ought to have, on the same considerations, and finally you predict that this rational agent will act to further its goals in the light of its beliefs. A little practical reasoning from the chosen set of beliefs and desires will in most instances yield a decision about what the agent ought to do; that is what you predict the agent will do."

=> seeing agents as intentional beings, acting according to beliefs and desires

The Intentional Stance (1987)

Daniell C. Dennett, Philosopher

http://www.glasbergen.com/

Belief-Desire-Intention Philosophy

Michael E. Bratman, Philosopher

- Bratman, (1987). Intention, plans, and practical reason

http://shc.stanford.edu/people/directory/michael-bratman

- Intention
  - mental attitude different from beliefs and desires
  - essential for theory of practical rationality

Belief-Desire-Intention Philosophy

Michael E. Bratman, Philosopher

Pro-attitudes play a motivational role: in concert with beliefs they can move us to act.

- intentions and desires are pro-attitudes, but motivational role is different
- desires influence future conduct, i.e., agent will be more inclined to act towards achieving its desires
- intentions control future conduct
- intentions involve commitment to action

http://shc.stanford.edu/people/directory/michael-bratman

4.

BDI Logics

Intention is Choice with Commitment
Cohen and Levesque (1990)

Desiderata for a theory of intention

- intentions pose problems for an agent; agent needs to determine how to achieve intentions
- intentions provide a "screen of admissability" for adopting other intentions
- agents track the success of their attempts to achieve their intentions; replan if earlier attempts fail
- ...

2006 Winner of IFAAMAS Influential Paper Award

Little Nell Problem
McDermott (1982)

Heroine (Nell) tied to train tracks while train is approaching. Dudley has to save Nell. Dudley reasons: "If Nell is going to be mashed, I must remove her from the tracks." Makes plan to save Nell. Taking this plan into account, Dudley reasons: "Nell will not be mashed." But then no justification for the plan anymore. Dudley removes plan (and starts over...).

Intention is Choice with Commitment
Cohen and Levesque (1990)

- goal: chosen desire
- by construction, goals are consistent
- persistent goal: goal that is kept as long as certain conditions hold
  - not give up until satisfied or impossible to achieve
  - commitment
- intention: kind of persistent goal

Modeling Rational Agents within a BDI-Architecture

Rao and Georgeff (1991)

- intention as first-class citizen in the logic, on a par with beliefs and goals
- goal: chosen desire
- chosen desires are consistent
- intention: goals that the agent committed to

2007 Winner of IFAAMAS Influential Paper Award

Summary

- goal: chosen desire
- goals are consistent
- intention: goals that the agent committed to

5.

Proactive Agents

Intelligent Agents: Theory & Practice
Wooldridge and Jennings (1995)

Weak agency

- autonomy
- social ability
- reactivity
- proactivity: able to exhibit goal-directed behavior by taking the initiative

http://www.csc.liv.ac.uk/~mjw/

http://users.ecs.soton.ac.uk/nrj/

Agent-Oriented Software Engineering
Jennings (1999)

An agent is an encapsulated computer system that is situated in some environment and that is capable of flexibel autonomous action in that environment in order to meet its design objectives.

Explicit Representation of Goals?

- standard OO program also exhibits goal-directed behavior (goal = postcondition): implicit goals
- need something more if change is the norm, not the exception: explicit representation of goals
  - plan failure decoupled from goal failure (Winikoff et al., 2002)
  - reasoning about goals (e.g., conflicts, goal adoption)

Conventional computer software is "task-oriented" rather than "goal-oriented"; that is, each task (or subroutine) is executed without any memory of why it is being executed. This means that the system cannot automatically recover from failures [...].
(Georgeff et al., 1999)

Motivation Frameworks
Hawes (2011)

- A system requires explicit goal representations which it can reason about.
- A system requires at least one process capable of creating goals.
- A system requires a process capable of collecting goals and then selecting which ones should be acted upon.
- A system requires a process which can generate goal directed behaviour from a collection goal and the available resources.

6.

Examples in GOAL

Wumpus

Elevator

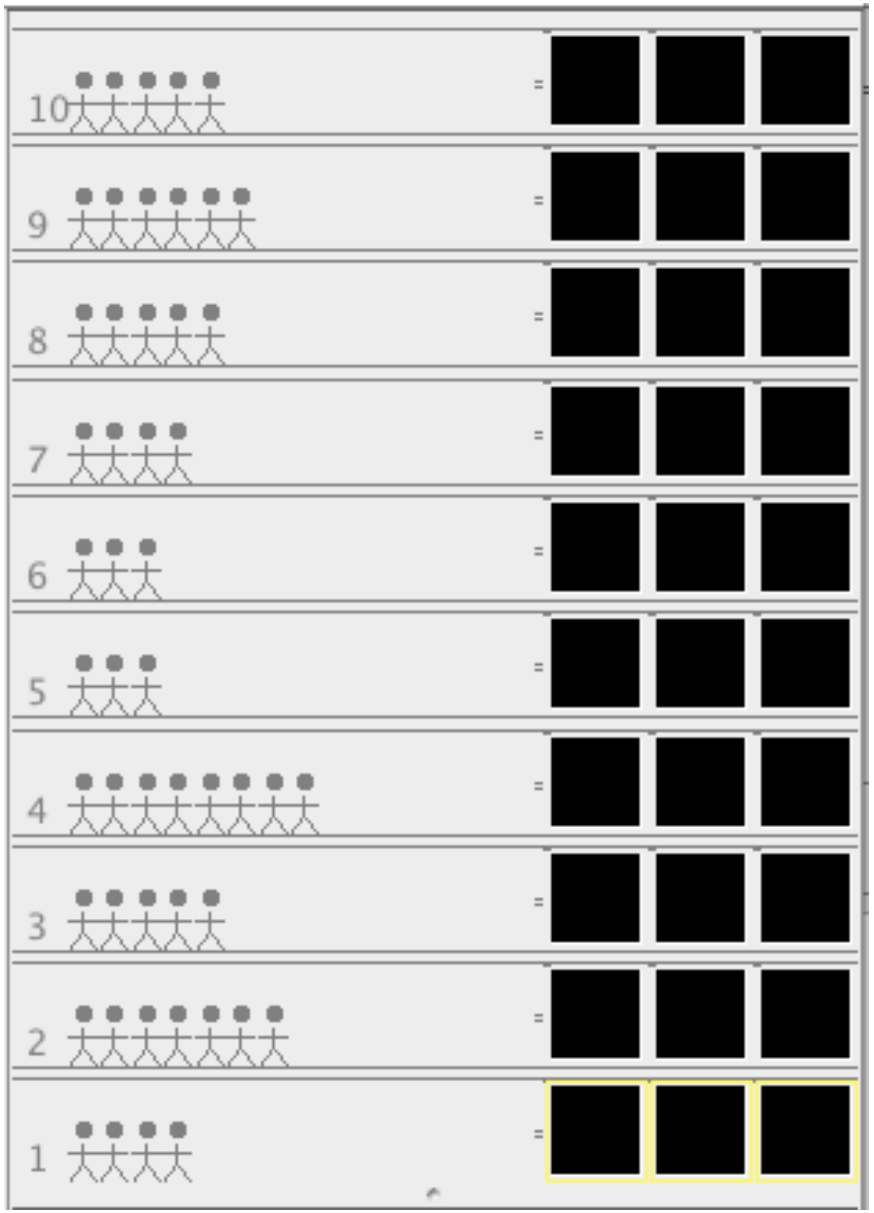

Goals:
at(Floor)
dir(Direction)

Tarzan & Jane

Goals:
meeting(date(1, 02, 2010),
time(12,00),
duration(1,00),
[jane, tarzan]).

http://www.allmoviephoto.com/photo/jane_tarzan_tarzan_001.html

GOAL STUDENT PROJECT

Goals:
has(Flag)
at(Location)
...

Summary

- BDI philosophy
  - desires and intentions as motivational attitudes
  - intentions control future conduct
- BDI logics
  - goals as consistent set of chosen desires
  - intention as goals that agent commits to
- Explicit representation of goals in agent systems
  - enables reasoning about goals