CPSC 589

Seminar in Computer Science

Agents based on BDI

Chang-Hyun Jo
Department of Computer Science
California State University Fullerton
jo@ecs.fullerton.edu

Contents

BDI Agents

- The BDI paradigm provides a strong notion of agency [Kinny et al 1996].
  - Agents are viewed as having certain mental attitudes, Belief, Desire and Intentions.
    - Informational, motivational and deliberative states

BDI Agents

- Belief-Desire-Intention (BDI) agents
  - Belief
    - Informative component of system state
    - Variables, DBs, a set of logical expressions, DSs
  - Desire
    - Motivational state of the system
  - Intention
    - Deliberative component of the system

BDI Agents

- Belief-Desire-Intention (BDI) architecture draws its inspiration from the philosophical theories of Bratman [Bratman 1987]
  - argues that intentions play a significant and distinct role in practical reasoning and cannot be reduced to beliefs and desires
- The 1st logical formalization of intentions and the notion of commitment is provided by Cohen and Levesque (1990).

BDI Agents

- Belief-Desire-Intention (BDI) formalization
  - Representation theory by Konolige and Pollack (1993)
  - Singh and Asher (1990)

BDI Agents

- Belief-Desire-Intention (BDI) agents
  - Theoretical foundations of BDI agents
    - Quantitative decision-theoretic perspective
    - Symbolic reasoning perspective
  - Implementations of BDI agents
  - Building of large-scale applications based on BDI agents

BDI Agents

- Belief-Desire-Intention (BDI) agents
  - Theoretical foundations of BDI agents
    - Bratman et al. 1987, Rao and Georgeff 1991, Shoham 1993
    - Cohen and Levesque 1990, Jennings 1992, Kinny et al. 1994
  - Implementations of BDI agents
    - Burmeister and Sundermeyer 1992, Georgeff and Lansky 1986, Shoham 1993
  - Building of large-scale applications based on BDI agents
    - Ingrand et al. 1992, Rao et al. 1992, OASIS (agent-oriented air-traffic management system tested at Sydney airport)

Abstract Architecture [Rao and Georgeff 1995]

- Three dynamic data structures
  - Belief
  - Desire
  - Intentions
- Input queue of events
  - External events
  - Internal events

BDI-Interpreter [Rao and Georgeff 1995]

```
initialize-state();
repeat
   options := option-generator(event-queue);
   selected-options := deliberate(options);
   update-intentions(selected-options);
   execute();
   get-new-external-events();
   drop-successful-attitudes();
   drop-impossible-attitudes();
end repeat
```


Abstract Architecture

- Idealization capturing the theory including the various components of practical reasoning [Bratman 1987]
  - Option generation, deliberation, execution, intention handling
- But, not a practical system for rational reasoning

Abstract Architecture

- The architecture is based on a logically closed set of beliefs, desires, and intentions.
- No indication of how the option generator and deliberation procedures can be made sufficiently fast to satisfy the real-time demands placed upon the system

More Practical Systems

- Procedural Reasoning Systems (PRS) [Georgeff and Lansky 1986] [Ingrand et al. 1992]
  - Agent-oriented systems based on the BDI architecture
- dMARS (distributed MultiAgent Reasoning System)
  - Its successor

Applications

- OASIS
  - An air-traffic management system
  - Related to the theoretical formalism and the abstract architecture
  - aircraft agent, sequencer, wind modeller, coordinator, trajectory checker

Applications

- PRS and dMARS
  - Space shuttle diagnosis [Ingrand et al. 1992]
  - Telecommunications network management [Ingrand et al. 1992]
  - Air-combat modelling [Rao et al. 1992]
  - Business process mangement [??]
  - Agent-oriented approach is useful for building complex distributed systems involving resource-bounded decision-making.

Essential Features

- The ability to construct plans that can react to specific situations, can be invoked based on their purpose, and are sensitive to the context of their invocation facilitates modular and incremental development.
- The balance between reactive and goal-directed behavior is achieved by committing to plans and periodically reconsidering such committed plans.

Essential Features

- The high-level representational and programming language has meant that end-users can encode their knowledge directly in terms of basic mental attitude without needing to master the programming constructs of a low-level language.
  - Greater flexibility, shorter development cycles [Rao et al 1992].

BDI Agent-based Modelling Technique [Kinny et al. 1996]

- There are not many examples of the successful application of agent system technologies on a significant scale.
  - Due to the absence of mature languages, software tools, and methodologies
  - OO methodologies are not directly applicable to agent systems.
    - Agents are more complex than typical objects, both in their internal structure and in the behaviors they exhibit.

BDI Agent-based Modelling Technique [Kinny et al. 1996]

- Explore how existing OO modeling techniques can be extended to apply to BDI agent systems
  - Kinny and Georgeff (1995)
  - Kinny et al. (1995)

Agent-Oriented Methodology

- OO methodologies [Booch 1994, Rumbaugh et al. 1991]
  - Object model
    - Information about objects
  - Dynamic model
    - State transitions, events, interactions that characterize system behavior
  - Functional model
    - Flow of data

Agent-Oriented Methodology

- In specifying an agent system:
  - External viewpoint
    - The system is decomposed into agents, modeled as complex objects characterized by their purpose, their responsibilities, the services they perform, the information they require and maintain, and their external interactions.
  - Internal viewpoint
    - The elements required by a particular agent architecture must be modeled for each agent.

Agents from the External Viewpoint

- Kinny et al. (1996) advocate the decomposition of a system based on the key role in an application.
- The identification of roles and their relationships guides the specification of an agent class hierarchy; agents are particular instances of these classes.

Agents from the External Viewpoint

- Analysis of the responsibilities of each agent class leads to the identification of the services provided and used by an agent, and hence its external interactions.
- Consideration of issues such as the creation and duration of roles and their interactions determines control relationship between agent classes.

Agents from the External Viewpoint

- Agent modeling from the external view point is captured in two models:
  - Agent Model
  - Interaction Model

Agents from the External Viewpoint

- An Agent Model describes the hierarchical relationship among different abstract and concrete agent classes, and identifies the agent instances which may exist within the system, their multiplicity, and when they come into existence.

Agents from the External Viewpoint

- An Interaction Model describes the responsibilities of an agent class, the services it provides, associated interactions, and control relationships between agent classes.
  - This includes the syntax and semantics of messages used for inter-agent communication and communication between agents and other system components, such as user interfaces.

Agents from the External Viewpoint

4 Steps to construct 2 Models

1. Identify the roles of the application domain.
   - Several dimensions in which such an analysis can be undertaken; roles can be organizational or functional, they can be directly related to the application, or required by the system implementation.
   - Identify the lifetime of each role.
   - Elaborate an agent class hierarchy*.
     - *criticism
   - The initial definition of agent classes should be quite abstract, not assuming any particular granularity of agency.

Agents from the External Viewpoint

4 Steps to construct 2 Models

2. For each role, identify its associated responsibilities, and the services provided and used to fulfill those responsibilities.
   - Services may include interactions with the external environment or human users.
     - Decompose agent classes to the service level*
       - * criticism

Agents from the External Viewpoint

4 Steps to construct 2 Models

3. For each service, identify the interactions associated with the provision of the service*, the performatives (speech acts) required for those interactions, and their information content.
   - Identify events and conditions to be noticed, actions to be performed, and other information requirements.
   - Determine the control relationship between agents.
   - At this point the internal modeling* of each agent class can be performed.
   - * criticism

Agents from the External Viewpoint

4 Steps to construct 2 Models

4. Refine the agent hierarchy*. (* criticism)
   - Where there is commonality of information or services between agent classes, consider introducing a new agent class, which existing agent classes can specialize, to encapsulate what is common.
   - Compose agent classes, via inheritance or aggregation, guided by commonality of lifetime, information and interfaces, and similarity of services.
   - Introduce concrete agent classes, taking into account implementation dependent considerations of performance, communication costs and latencies, fault-tolerance requirements, etc.
   - Refine the control relationships.
   - Finally, based upon considerations of lifetime and multiplicity, introduce agent instances.

Agents from the External Viewpoint

- Roles, responsibilities, and services are just descriptions of purposeful behaviors at different levels of abstraction.
  - Roles can be seen as sets of responsibilities, and responsibilities as sets of services.
  - Services are those activities that it is not natural to decompose further, in terms of the identity of the performer*. (* criticism)
  - The roles initially identified serve as a starting point for the analysis.
  - Once roles have been decomposed to the level of services and internal modeling performed, a fine-grained model of agency has been produced.

Agents from the External Viewpoint

- Subjects of ongoing research
  - Service relationships and interactions
  - Information contents and linguistic intent through to protocols for coordination and negotiation [See the reference lists.]
  - Interaction models capturing control relationships between agents, such as responsibilities for agent creation and deletion, delegation, and team formation
  - * criticism (dynamic configuration)

Agents from the Internal Viewpoint

- In the BDI architecture, an agent can be completely specified by the events that it can perceive, the actions it may perform, the belief it may hold, the goal it may adopt, and the plans that give rise to its intentions.

Agents from the Internal Viewpoint

- A Belief Model describes the information about the environment and internal state that an agent of that class may hold, and the actions it may perform.
  - The possible beliefs of an agent and their properties, such as whether or not they may change over time, are described by a belief set.
  - In addition, one or more belief states – particular instances of the belief set – may be defined and used to specify an agent's initial mental state.

Agents from the Internal Viewpoint

- A Goal Model describes the goals that an agent may possibly adopt, and the events to which it can respond.
  - It consists of a goal set which specifies the goal and event domain and one or more goal states – sets of ground goals – used to specify an agent's initial mental state.

Agents from the Internal Viewpoint

- A Plan Model describes the plans that an agent may possibly employ to achieve its goals.
  - It consists of a plan set which describes the properties and control structure of individual plans.

Agents from the Internal Viewpoint

- Implicit in this characterization are the execution properties of the architecture which determine how, exactly, events and goals give rise to intentions, and intentions lead to action and revision of beliefs and goals.

Agents from the Internal Viewpoint

- These properties are responsible for ensuring that beliefs, goals, and intentions evolve rationally [Kinny 1993].
  - The architecture ensures that events are responded to in a timely manner, beliefs are maintained consistently, and that plan selection and execution proceeds in a manner which reflects certain notions of rational commitment.

Agents from the Internal Viewpoint

How to get models

- Begins from the services provided by an agent* and the associated events and interactions.
  - These define its purpose, and determine the top-level goals that the agent must be able to achieve.
  - Analysis of the goals and their further breakdown into sub-goals leads naturally to the identification of different means, i.e., plans, by which a goal can be achieved.
- (* criticism)

Agents from the Internal Viewpoint

How to get models

- The appropriateness of a given plan, and the manner in which a plan is carried out, will in general depend upon the agent's beliefs about the state of the environment and possibly other information available to the agent, i.e., the agent's belief context.
  - A context is represented in terms of various data entities and their relationships.
  - Analysis of contexts results in the elaboration of the beliefs of an agent.

Agents from the Internal Viewpoint

Two steps for internal modeling methodology

1. Analyze the means of achieving the goals.
   - For each goal, analyze the different contexts in which the goal has to be achieved.
   - For each of these contexts, decompose each goal into activities, represented by sub-goals, and actions.
   - Analyze in what order and under what conditions these activities and actions need to be performed, how failure should be dealt with, and generate a plan to achieve the goal.
   - Repeat the analysis for sub-goals.

Agents from the Internal Viewpoint

Two steps for internal modeling methodology*

2. Build the beliefs of the system.
   - Analyze the various contexts, and the conditions that control the execution of activities and actions, and decompose them into component beliefs.
   - Analyze the input and output data requirements for each sub-goal in a plan and make sure that this information is available either as beliefs or as outputs from prior sub-goals in the plan.
   - (* evaluation)

Agents from the Internal Viewpoint

- These steps are iterated as the models which capture the results of analysis are progressively elaborated, revised, and refined.
- Refinement of the internal models feeds back to the external models.
  - Building the plans and beliefs of an agent class clarifies the information requirements of services, particularly with respect to monitoring and notification.

Agents from the Internal Viewpoint

- Key abstractions to manage complexity
  - Roles, responsibilities, services, goals
- Goal-oriented design
  - Stable, robust, and modular design
  - Goals compared to behaviors or plans, are more stable in any application domain
  - Correctly identifying goals leads to a more robust system design, where changes to behaviors can be accommodated as new ways of achieving the same goal.

Agents from the Internal Viewpoint

- The context-sensitivity of plans provides modularity and compositionality.
  - Plans for new contexts may be added without changing existing plans for the same goal.
  - This results in an extensible design that can cope with frequent changes and special cases, and permits incremental development and testing.

Agent-Oriented Modeling

- Case study [Kinny 1996]
  - Air-traffic management system

Modeling Technique

- Kinny (1996)

Modeling and Design of Multi-Agent Systems

- Modeling and design of multi-agent systems [Kinny and Georgeff 1996]
  - The BDI architecture based framework

References

- Acknowledgement
  - These notes are summarized mainly from the following references.

  - Rao, Anand S. and Georgeff, Michael P., BDI Agents: From Theory to Practice, Proc. of the 1st International Conference on Multi-Agent Systems (ICMAS-95), San Francisco, USA, June, 1995.
  - Kinny, D. Georgeff, M., and Rao, A., A Methodology and Modelling Technique for Systems of BDI Agents, In Agents Breaking Away: The 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'96, Springer, Lecture Notes in AI 1038, 56-71, 1996.
  - Kinny, D. and Georgeff, M., Modelling and Design of Multi-Agent Systems, Intelligent Agents III: The 3rd International Workshop on Agent Theories, Architectures, and Languages (ATAL-96), Springer, Lecture Notes in AI, 1996.

References

- Kinny, D. and Georgeff, M. Modelling techniques for BDI agents systems, Technical Report 54, Australian AI Institute, Melbourne, Australia, 1995.
- Kinny, D., Georgeff, M., and Rao A., A design methodology for BDI agent systems, Technical Report 55, Australian AI Institute, Melbourne, Australia, 1995.
- Bratman, M. E., Intention, Plans, and Practical Reason, *Harvard University Press*, 1987. (also available from CSLI Publications, 1999)