

Chapter 1

JADEX: A BDI REASONING ENGINE

Alexander Pokahr,¹ Lars Braubach,¹ and Winfried Lamersdorf¹

¹*University of Hamburg
Distributed Systems and Information Systems
22527 Hamburg, Germany
{pokahr|braubach|lamersd}@informatik.uni-hamburg.de*

Abstract This chapter presents Jadex, a software framework for the creation of goal oriented agents following the belief-desire-intention (BDI) model. The Jadex project aims to make the development of agent based systems as easy as possible without sacrificing the expressive power of the agent paradigm. The objective is to build up a rational agent layer that sits on top of a middleware agent infrastructure and allows for intelligent agent construction using sound software engineering foundations. Fostering a smooth transition from traditional distributed systems to the development of multi-agent systems, well established object-oriented concepts and technologies such as Java and XML are employed wherever applicable. Moreover, the Jadex reasoning engine tries to overcome traditional limitations of BDI systems by introducing explicit goals. This allows goal deliberation mechanisms being realized and additionally facilitates application development by making results from goal oriented analysis and design easily transferable to the implementation layer. The system is freely available under LGPL license and provides extensive documentation as well as illustrative example applications.

Keywords: BDI agents, FIPA standard, object-oriented software engineering, explicit goals.

1.1 Motivation

Today, a numerousness of different agent platforms is available for developing multi-agent applications [13]. Nevertheless, most of these platforms are developed with a specific technological focus such as the

cognitive or infrastructural architecture. Hence, not all aspects of agent technology are covered equally well. General applicability of an agent platform for a great variety of domains demands that at least three categories of requirements are considered: openness, middleware and rationality. Openness is closely related to the vision of interconnected networks of originally unrelated applications whereas middleware aspects emphasize traditional software engineering concerns such as service management, security and persistency aspects. Rationality, in turn, focusses on the agent's internal decision making process and mostly tries to map this process from a natural archetype such as insects or humans.

According to these aspects the existing platforms can be classified into two almost distinct groups. On the one hand, FIPA-compliant platforms mainly address openness and middleware issues by realizing the FIPA communication resp. platform standards [19]. On the other hand, reasoning-centered platforms exist, that focus on the behaviour model of a single agent trying to achieve rationality and goal-directedness. This gap between middleware and reasoning-centered systems is one main motivation for the realization of the Jadex BDI (Belief-Desire-Intention) reasoning engine [5, 18], which aims to bring together both research strands.

Besides this overall objective to support both FIPA compliance and BDI reasoning, the design of the system is driven by two main factors. On the one hand, the development of the reasoning engine is accompanied by an ongoing effort of enhancing the BDI architecture in general. The system addresses shortcomings of earlier BDI agent systems, e.g. providing an explicit representation of goals and a systematic way for the integration of goal deliberation mechanisms. On the other hand, the system respects the current state of the art regarding mainstream (OO) software engineering, and is designed to be used not only by AI experts, but by the ordinarily skilled software developer. Therefore, agent development is based on established techniques such as Java and XML, and is further supported by software engineering aspects, such as reusable modules and development tools.

1.2 Architecture

This section presents the architectural underpinnings of the Jadex system. It starts with a short review of the BDI model. Subsequently, an overview of the architecture of Jadex is presented. The basic concepts – beliefs, goals, and plans – of the system are introduced by highlighting their main characteristics and differences to other BDI agent systems.

Finally, the execution model is shortly sketched, showing how the components of the system interoperate.

1.2.1 BDI Model

The BDI model was initially conceived by Bratman as a theory of human practical reasoning [3]. Its success is based on its simplicity reducing the explanation framework for complex human behavior to the *motivational stance* [9]. In this model, causes for actions are only related to desires ignoring other facets of cognition such as emotions. Another strength of the BDI model is the consistent usage of folk psychological notions that closely correspond to the way people communicate about human behavior [14]. The BDI theory of Rao and Georgeff [20] defines beliefs, desires, and intentions as mental attitudes represented as possible world states. The intentions of an agent are subsets of the beliefs and desires, i.e., an agent acts towards some of the world states it desires to be true and believes to be possible. To be computationally tractable Rao and Georgeff also proposed several simplifications to the theory, the most important one being that only beliefs are represented explicitly. Desires are reduced to events that are handled by predefined plan templates, and intentions are represented implicitly by the runtime stack of executed plans.

1.2.2 Concepts within Jadex

In Fig. 1.1 an overview of the abstract Jadex architecture is presented. Viewed from the outside, an agent is a black box, which receives and sends messages. As common in BDI systems, all kinds of events, such as incoming messages or goal events serve as input to the internal reaction and deliberation mechanism, which dispatches the events to plans selected from the plan library. The reaction and deliberation mechanism is the only global component of an agent. All other components are grouped into reusable modules called capabilities.

Beliefs. One objective of the Jadex project is the adoption of a software engineering perspective for describing agents. Therefore, Jadex does not enforce a logic-based representation of beliefs. Instead, an object oriented representation of beliefs is employed, where arbitrary objects are stored as named facts (called beliefs) or named sets of facts (called belief sets). Operations against the beliefbase can be issued in a descriptive set-oriented query language. Moreover, the beliefbase is not only a passive data store, but takes an active part in the agent's execution, by monitoring belief state conditions. Changes of beliefs may

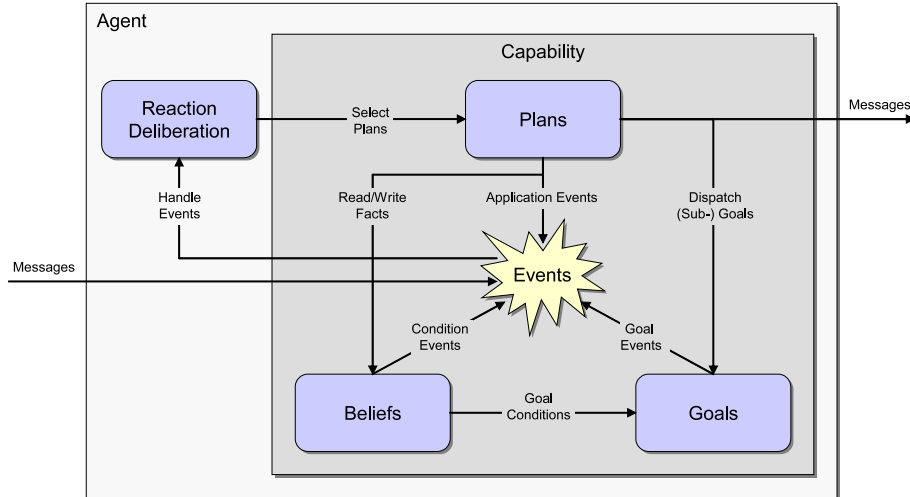


Figure 1.1. Jadex abstract architecture

therefore directly lead to actions such as events being generated or goals being created or dropped.

Goals. Unlike traditional BDI systems, which treat goals merely as a special kind of event, goals are a central concept in Jadex. Jadex follows the general idea that goals are concrete, momentary desires of an agent. For any goal it has, an agent will more or less directly engage into suitable actions, until it considers the goal as being reached, unreachable, or not desired any more. Unlike most other systems, Jadex does not assume that all adopted goals need to be consistent to each other. To distinguish between just adopted (i.e. desired) goals and actively pursued goals, a goal lifecycle is introduced which consists of the goal states *option*, *active*, and *suspended* (see Fig. 1.2). When a goal is adopted, it becomes an option that is added to the agent's desire structure. Application specific goal deliberation mechanisms are responsible for managing the state transitions of all adopted goals (i.e. deciding which goals are active and which are just options). In addition, some goals may only be valid in specific contexts determined by the agent's beliefs. When the context of a goal is invalid it will be suspended until the context is valid again.

Jadex supports four types of goals, which extend the general lifecycle and exhibit different behaviour with regard to their processing as explained below. A *perform* goal is directly related to the execution of actions. Therefore the goal is considered to be reached, when some ac-

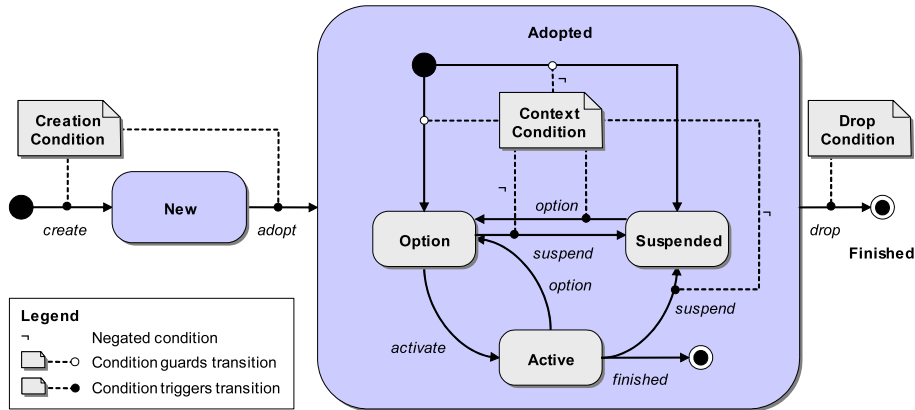


Figure 1.2. Goal lifecycle

tions have been executed, regardless of the outcome of these actions. An *achieve* goal is a goal in the traditional sense, which defines a desired outcome without specifying how to reach it. Agents may try several different alternative plans, to achieve a goal of this type. A *query* goal is similar to an achieve goal. Its outcome is not defined as a state of the world, but as some information the agent wants to know about. For goals of type *maintain*, an agent keeps track of the desired state, and will continuously execute appropriate plans to re-establish the maintained state whenever needed. More details about goal representation and processing in Jadex can be found in [7].

Plans. Plans represent the behavioural part of an agent and are composed of a head and a body part. The plan head specification is similar to other BDI systems and mainly specifies the circumstances under which a plan may be selected, e.g. by stating events or goals handled by the plan and preconditions for the execution of the plan. Additionally, in the plan head a context condition can be stated that must be true for the plan to continue executing. The plan body provides a predefined course of action, given in a procedural language. This course of action is to be executed by the agent, when the plan is selected for execution, and may contain basic actions provided by the system, such as sending messages, manipulating beliefs, or creating subgoals.

Capabilities. Capabilities, introduced in [8], represent a grouping mechanism for the elements of a BDI agent, such as beliefs, goals, plans, and events. In this way, closely related elements can be put together into

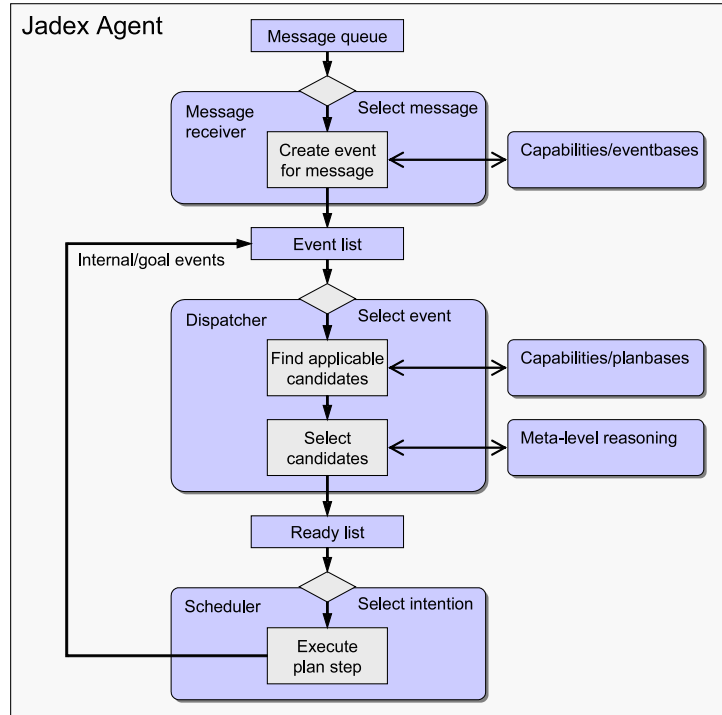


Figure 1.3. Jadex execution model

a reusable module, which encapsulates a certain functionality (e.g. for interaction with a FIPA directory facilitator). The enclosing capability of an element represents its scope, and an element only has access to elements of the same scope (e.g. a plan may only access beliefs or handle goals or events of the same capability). To connect different capabilities, flexible import / export mechanisms can be used that define the external interface of the capability (e.g. beliefs or goals visible to the outside).

1.2.3 Execution Model

This sections shows the operation of the reaction and deliberation component, given the Jadex BDI concepts as described earlier. All of the required functionality is assigned to cleanly separated components, which will be explained in turn. Incoming messages are placed in the agent's global message queue by the underlying agent platform. Before the message can be forwarded the system, it has to be assigned to one or more capabilities, which are able to handle the message. If the message

belongs to an ongoing conversation (identified by a conversation-id as part of the message), an event for the incoming message is created in the capability executing the conversation. Otherwise, suitable capabilities have to be found, which is done by matching the message against event templates defined in the eventbase of each capability. These matching templates are then used to create appropriate events in the scope of the capabilities. In either case, the created events are subsequently added to the agent's global event list.

The dispatcher is responsible for selecting applicable plans for the events from the event list. This is done in two steps: First, a list of applicable plans is generated by matching the event against the plan heads as defined in the planbases of each capability. Only those capabilities have to be considered, where the event is defined. The second step is to select a subset of the applicable plans for execution. Regarding this step several important questions arise, such as if all of the applicable plans should be executed concurrently, or if the event is posted to another plan if the first plan fails [8]. The decision, which plan to execute is called meta-level reasoning and may be as simple as selecting the first plan from the list, or as complicated as finding and executing meta-plans for the decision. In Jadex an event provides flexible settings to influence processing on the basis of individual events. As a default, messages are posted to only a single plan, while for goals, many plans are executed sequentially until the goal is finished. Other internal events are posted to all plans at once as they are considered only as a change notification and no result is expected. After plans have been selected, they are placed in the ready list, waiting for execution.

The execution of plans is performed by the scheduler, which selects the plans from the ready list. Plans are executed step-by-step, whereby the length of plan step depends on the context. A plan step is defined to be any sequence of basic actions, until the plan waits explicitly or affects the internal state of the agent. After each plan step the state of the agent is properly updated (e.g. goals are checked for achievement or failure), maybe leading to new internal or goal events. The dynamic plan step length assures atomic agent state changes without limiting plan steps to executing only a single predefined action.

1.3 Language

Jadex is neither based on a new agent programming language nor does it employ or revise an existing one. Instead, a hybrid approach is chosen, distinguishing explicitly between the language used for static agent type specification and the language for defining the dynamic agent

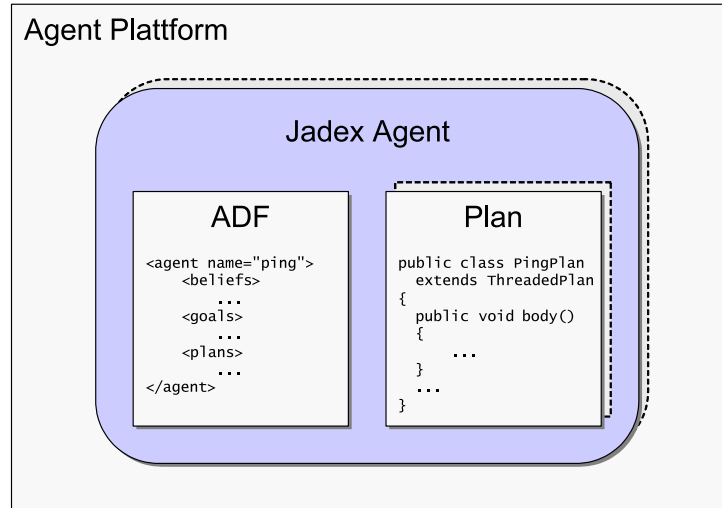


Figure 1.4. Jadex agent

behaviour. According to this distinction, a Jadex agent consists of two components: An agent definition file (ADF) for the specification of inter alia beliefs, goals and plans as well as their initial values and on the other hand procedural plan code (see Fig. 1.4). For defining ADFs an XML language is used that follows the Jadex BDI metamodel specified in XML Schema. In contrast, the procedural part of plans (the plan bodies) are realized in an ordinary programming language (Java) and have access to the BDI facilities of an agent through an application program interface (API).

1.3.1 Specifications and Syntactical Aspects

The Jadex BDI meta model defined in XML Schema is very extensive and hence cannot be presented completely in this paper (for a complete introduction see [17]). Generally, the corresponding language was specified with two design principles in mind. First design objective is the support for strong typing and explicit representation of all kinds of elements, be it beliefs, goals or events. In consequence, this requires users to write detailed ADFs, but in return allows for the realization of more rigorous consistency checking of agent models. Additionally, at runtime certain kinds of failures can be discovered more easily, e.g. the attempt of storing a fact value in an undefined belief can be immediately reported.


```

01: select_expression ::= "SELECT" ("ALL" | "ANY" | "IOTA")?
02: (
03:   (expression "FROM" ("$" identifier "IN" expression) ("," "$" identifier "IN" expression)* )
04:   | ("$" identifier "FROM" expression)
05: )
06: ("WHERE" expression)?
07: ("ORDER" "BY" expression ("ASC" | "DESC")?)?

```

Figure 1.5. OQL syntax in EBNF

Second design objective regards increasing the expressive power of XML for the following purposes: The arbitrary complex creation of objects (e.g. values within beliefs or parameters), the description of boolean conditions (e.g. when a certain goal should be dropped) and the construction of queries (e.g. for retrieving values from the beliefbase). To achieve this, an embedded language is used for specifying parts of the agent model, not easily represented in XML. This expression language is used for the interpretation of text elements within an ADF and has been designed to fully comply with the syntax of Java expressions (right hand side of assignments) extended with a subset of OQL (object query language) instructions. The syntax of the OQL extension is depicted in Fig. 1.5 in EBNF notation. It allows for query statements being created in the well-known *select-from-where* form, whereby it can be additionally specified if exactly one (iota), the first satisfying (any) or all satisfying results are expected (line 1). In the *from* clause (lines 03-04) it is specified from which object set (line 4) or joined sets (line 3) results are generated, whereby the identifiers define variables, which iterate over the object sets specified as arbitrary expressions. These iterated values are checked against the boolean *where* condition (line 6) and can possibly be ordered (line 7).

In the following the essential BDI concepts as presented in Section 1.2.2 will be taken on and their realization on language level will be detailed. These concepts are specified as part of an agent or capability description in the same manner. In Fig. 1.6 (left hand side) the allowed attributes and subtags of the agent tag are shown. Each agent type is identified by a name and package declaration and can be provided with a description text. In addition, the corresponding agent class and runtime properties can be set. For most cases the default values are sufficient and need not to be modified. It can be seen that besides the subtags for the core BDI concepts (beliefs, goals, plans and events which are explained below) several other elements can be declared. Most of these elements (languages, ontologies, servicedescriptions and agentdescriptions) are FIPA related and facilitate agent communication resp. the interaction with yellow page services. The remaining elements (im-

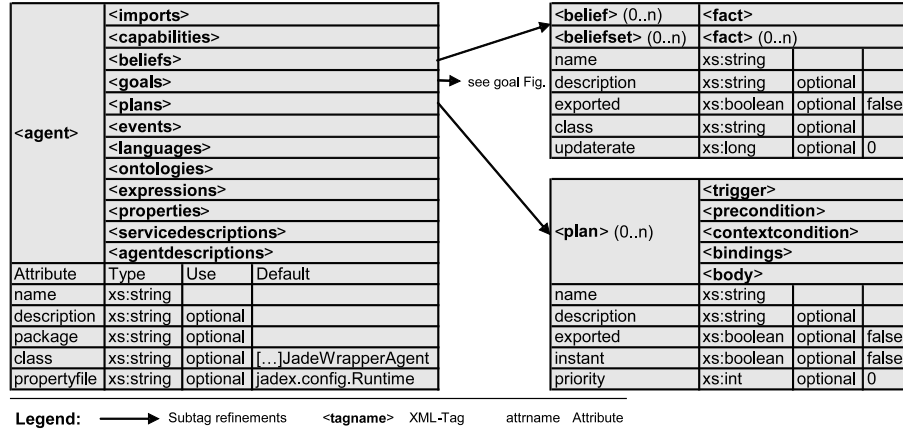


Figure 1.6. Agent metamodel specification fragment (XML-schema)

ports, expressions, properties) are of less importance and only serve for convenience and agent configuration purposes.

In Jadex beliefs are represented in an object-oriented way allowing arbitrary Java objects being stored as facts. Like all scoped elements, beliefs and belief sets can be supplied with a name, a description text and an exported flag. Exporting an element makes it accessible from the outer scope (resp. a capability or an agent) and is turned off by default. For beliefs and belief sets the Java class for facts must be defined. Besides the type-relevant information, initial fact data can also be supplied for configuring an agents mental state at creation time. The value of a fact has to be stated in the expression language and can be declared as static or dynamic, whereby dynamic facts are useful e.g. for representing values continuously sensed from an environment or time-relevant aspects. Recalculation of such dynamic facts occurs on access and additionally in fixed time intervals (using the update rate). At runtime, beliefs and belief sets are accessible from within plans via operations on the belief base and additionally by issuing OQL-like queries.

As described earlier in Jadex between four different goal types is distinguished (perform, achieve, maintain and query). All these goal types are based on the generic life cycle and hence exhibit many common properties that are summarized in an abstract base goal type (see Fig. 1.7). According to the life cycle, creation, drop and context conditions can be specified as boolean expressions. Customization of goal types can be further achieved by defining named in-, out- and inout-parameters that are used to transfer information between a goal's originator and its process-

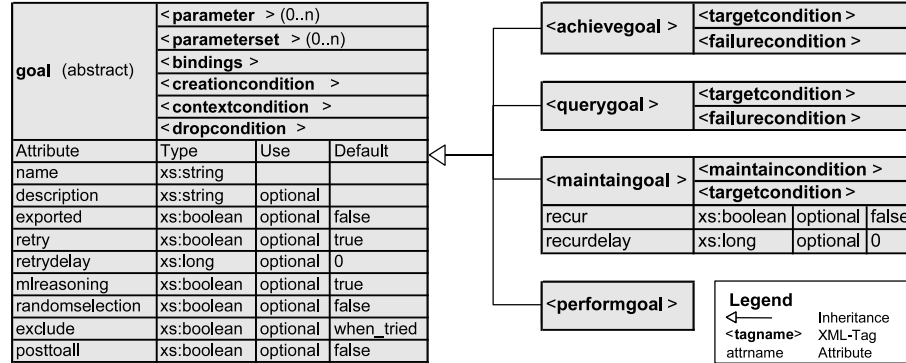


Figure 1.7. Goal metamodel specification (XML-schema)

ing plans. Additionally, binding parameters can be used for generating one goal instance for every possible binding. The runtime processing of goals can be refined using the various BDI-flags, which inter alia control if a goal is retried when a plan fails (retry), if meta-level reasoning is used (mreasoning) and if applicable plans are tried sequentially or in parallel (posttoall). A complete explanation can be found in [17]. From this abstract goal type all concrete types are derived. The simplest one being the perform goal used for executing (possibly repeatedly) certain actions, which does not require extra specification data. An achieve goal extends this abstract goal type and adds support for the specification of a target and a failure condition. The target condition is used for describing the world state this goal seeks to bring about as a boolean expression. Similar, a boolean failure condition has the purpose to abort goal processing in case its achievement has become impossible. The query goal provides the same kind of conditions, but exhibits a slightly different behaviour in that it is used for information retrieval purposes. Most complex behaviour is exposed by the maintain goal type, which is used to monitor a specific world state (maintain condition) and automatically tries to reestablish this state whenever it becomes invalid. A boolean target condition can be used to refine the state that is tried to be restored. Maintain goals are per se persistent and would remain inactive when the monitored state is violated and could not be reestablished at once. Hence, a maintain goal can be configured to retry re-establishment in certain time intervals (recur and recurdelay). In addition to the specification of the four types of goals, possibly parametrized initial goals can be declared that will be created when the agent is born. At runtime, goal instances can be created from within plans by referring to their

type name. Typically, some parameter values need to be supplied before a goal can be dispatched as top-level goal or as subgoal of the current plan.

The declaration of plans in Jadex is very similar to other BDI systems and requires the specification of the plan heads describing the circumstances under which a plan is applicable in the ADF. As plan trigger, events (including messages and goals) and a belief state condition (for data driven plans) can be provided; the pre- and context condition of a plan can be specified as boolean expressions. To facilitate goal achievement with plans, it is sometimes advantageous to create several different parametrized plan instances of a plan type and try them one after another until a plan succeeds. For this purpose binding parameters can be specified and used for plan configuration. Furthermore, the selection of which plan is executed in response to an occurring trigger can be adjusted by setting a priority value. As part of the initial mental state of an agent it can be further declared whether a plan is instantiated when the agent is created (using the instant flag).

The plan body needs to be supplied as expression for the creation of a suitable plan instance. Currently two different types of plan bodies (threaded, and non-threaded) are supported, which both require a Java class to be implemented. Non-threaded plan bodies have several disadvantages compared to the threaded versions, but nonetheless make sense in mobile scenarios as agent migration is provided. In Fig. 1.8 the skeleton of an application plan is depicted. Mandatory is only the extension of a corresponding framework class (`ThreadedPlan`) and the implementation of the abstract `body()` method, in which the domain-specific plan behaviour can be placed. In addition to the body method, three other methods exist that optionally can be implemented. These methods are called when plan processing has finished according to the plans final state. The `passed()` method is called when the body method runs through, whereas the `failure()` method is invoked when an uncaught exception is thrown within the `body()` method. Finally, the `aborted()` method is called, when plan processing was interrupted from outside. Two different abort cases can be distinguished, either when the corresponding goal succeeds before the plan is finished or when the plans root goal is dropped.

1.3.2 Software Engineering Issues

The overall goal of the Jadex project is to provide a sophisticated reasoning engine allowing to develop arbitrary complex intelligent agents. Therefore, while trying to be as easily useable as possible, the system

```
01: /** Plan skeleton for an application plan. */
02: public class SomePlan extends jadex.runtime.ThreadedPlan {
03:
04:     public void body() {
05:         // Plan code.
06:     }
07:
08:     public void passed() {
09:         // Optional cleanup code in case of a plan success.
10:     }
11:     public void failure() {
12:         // Optional cleanup code in case of a plan failure.
13:     }
14:     public void aborted() {
15:         // Optional cleanup code in case the plan is aborted.
16:     }
17: }
```

Figure 1.8. Plan skeleton of a ThreadedPlan

does not sacrifice expressiveness for simplicity. Nonetheless, software engineering issues play an important role in the design of the system.

As stated earlier, a primary goal of the project is to facilitate a smooth transition from mainstream object-oriented software development to an agent oriented approach. This is achieved by resorting to established techniques wherever possible. E.g., the system builds on Java and XML, therefore the developer does not have to learn a new language. Another great advantage is that the developer can continue to operate in a familiar environment. As the agent developer only has to create Java and XML files, existing development environments such as Eclipse¹ or IntelliJ IDEA² can be used to develop Jadex agents. In recent editions of these environments, features such as on-the-fly checking and auto-completion not only apply to Java code but can also easily be adopted for XML ADF creation,³ therefore offering complete support for Jadex agent development.

Moreover, the system provides advanced software engineering features, such as reusability and consistency checking. The capability concept allows to encapsulate agent functionality into a reusable module while maintaining the abstraction level of BDI elements. The explicit specification and strong typing of beliefs, goals, etc. facilitates consistency checks to detect errors (e.g. spelling mistakes) as early as possible.

¹<http://www.eclipse.org/>

²<http://www.jetbrains.com/idea/>

³In eclipse this is realized by the XMLBuddy plugin (see <http://xmlbuddy.com/>).

1.3.3 Example

To further explain the syntax and semantics of the Jadex agent languages, in this section a simple example is provided. The example does only cover a small subset of the features of Jadex. Another example covering all different types of goals can be found elsewhere [7]. The example presented here is a fully functional agent, taken directly from the current Jadex distribution. Purpose of the agent is to establish given configurations in a blocksworld environment, where colored blocks are placed in stacks on top of a table. The example provides a graphical user interface, where the user can visually create custom block configurations. The configurations have to be established by the agent by moving the blocks. As only clear blocks (without other blocks on top) can be moved, the agent has to perform some ad-hoc planning. The implemented solution is very simple, creating the stacks bottom-to-top, moving all obstructing blocks to the table. Details of GUI handling are omitted in the following, as they are not required to understand the behaviour of the agent. First, the model as given in XML will be described. Thereafter, the two plan bodies used by the agent and written in Java will be presented.

The ADF of the agent is shown in Fig. 1.9, where tags (elements of the Jadex metamodel) are in boldface, and embedded expressions are in italics. The model starts with the declaration of the agent tag, specifying the name and package of the agent (line 1). The package is used as first place to resolve references to other files such as capabilities and Java classes. More packages and files can be explicitly specified in the imports section (lines 2-4). In this case the class `java.lang.Color` is imported, because it is used to represent the color of a block.

The beliefs of the agent are given in the beliefs section (lines 6-16). A belief “table” (lines 7-9) is used to represent the environment, which consists of a table on which blocks are located. As initial fact of the belief, an instance of the `Table` class (located in package `jadex.examples.blocksworld`) is created (line 8). The known blocks are collected in a belief set “blocks” (lines 10-15). A number of blocks (class `Block`) with different colors is initially created given by single fact items (lines 11, 12, 13, ...). Two blocks are created on top of the table, while the third block is located on the first block (referenced by `table.allBlocks[0]`).

The agent has three achieve goals, each with a name, parameters and a corresponding target condition (lines 18-34). The “clear” goal (lines 19-22) represents the goal to clear (i.e. remove blocks located on top) a block given in a parameter (line 20). The target condition (line 21) refers directly to the `isClear()` method of this block. The “stack” goal

```

01: <agent name="Blocksworld" package="jadex.examples.blocksworld">
02:   <imports>
03:     <import>java.awt.Color</import>
04:   </imports>
05:
06:   <beliefs>
07:     <belief name="table" class="Table">
08:       <fact>new Table()</fact>
09:     </belief>
10:     <beliefset name="blocks" class="Block">
11:       <fact>new Block(new Color(240, 16, 16), $beliefbase.table)</fact>
12:       <fact>new Block(new Color(16, 16, 240), $beliefbase.table)</fact>
13:       <fact>new Block(new Color(240, 240, 16), $beliefbase.table.allBlocks[0])</fact>
14:       ...
15:     </beliefset>
16:   </beliefs>
17:
18:   <goals>
19:     <achievegoal name="clear">
20:       <parameter name="block" class="Block" />
21:       <targetcondition>$goal.block.isClear()</targetcondition>
22:     </achievegoal>
23:     <achievegoal name="stack">
24:       <parameter name="block" class="Block" />
25:       <parameter name="target" class="Block" />
26:       <targetcondition>$goal.block.lower==$goal.target</targetcondition>
27:     </achievegoal>
28:     <achievegoal name="configure">
29:       <parameter name="configuration" class="Table" />
30:       <targetcondition>
31:         $beliefbase.table.configurationEquals($goal.configuration)
32:       </targetcondition>
33:     </achievegoal>
34:   </goals>
35:
36:   <plans>
37:     <plan name="stack">
38:       <body>new StackBlocksPlan($event.goal.block, $event.goal.target)</body>
39:       <trigger><goal ref="stack"/></trigger>
40:     </plan>
41:     <plan name="configure">
42:       <body>new ConfigureBlocksPlan($event.goal.configuration)</body>
43:       <trigger><goal ref="configure"/></trigger>
44:     </plan>
45:     <plan name="clear">
46:       <bindings>
47:         <binding name="upper">
48:           select $upper from $beliefbase.blocks where $upper.getLower()==$event.goal.block
49:         </binding>
50:       </bindings>
51:       <body>new StackBlocksPlan($upper, $beliefbase.table)</body>
52:       <trigger><goal ref="clear"/></trigger>
53:     </plan>
54:   </plans>
55: </agent>

```

Figure 1.9. Blocksworld agent model

```

01: package jadex.examples.blocksworld;
02: import jadex.runtime.*;
03:
04: /** Plan to stack one block on top of another target block. */
05: public class StackBlocksPlan extends ThreadedPlan {
06:     protected Block block;
07:     protected Block target;
08:
09:     public StackBlocksPaperPlan(Block block, Block target) {
10:         this.block = block;
11:         this.target = target;
12:     }
13:
14:     public void body() {
15:         IGoal clear = createGoal("clear");
16:         clear.getParameter("block").setValue(block);
17:         dispatchSubgoalAndWait(clear);
18:
19:         clear = createGoal("clear");
20:         clear.getParameter("block").setValue(target);
21:         dispatchSubgoalAndWait(clear);
22:
23:         block.stackOn(target);
24:     }
25: }

```

Figure 1.10. Java code for StackBlocksPlan

(lines 23-27) aims at placing a given block (line 24) on a target block (line 25). Achieving this goal means that the block below the first block is now equal to the target block as stated by the target condition (line 26). To establish a complete configuration of blocks on the table, the “configure” goal (lines 28-31) is used. The desired configuration is given as a parameter of type Table (line 29). The target condition (line 30-32) refers to the configurationEquals() method implemented in the Table class. No initial instances of these three goal types are defined in the model. The agent starts idle, waiting for goals to appear, which are created by the user through a GUI.

The goals are handled by the plans of the agent (lines 36-54). In this example there is one plan for each goal, although this kind of one to one mapping is not required. The plan head declarations of the first two plans “stack” (lines 37-40) and “configure” (lines 41-44) are straightforward. The trigger (lines 39 resp. 43) defines when the plan is applicable, in this case for goals of type “stack” and “configure”, respectively. The body (lines 38, 42) defines how the plan body object is instantiated. In both cases, the creation expression refers to parameters of the triggering goal to supply the arguments for the Java constructor (cf. Figs. 1.10, 1.11). The “clear” plan definition is more complex, as the body of the


```

01: package jadex.examples.blocksworld;
02: import jadex.runtime.*;
03:
04: /** Plan to to establish a given configuration of blocks. */
05: public class ConfigureBlocksPlan extends ThreadedPlan {
06:     protected Table table;
07:
08:     public ConfigureBlocksPaperPlan(Table table) {
09:         this.table = table;
10:     }
11:
12:     public void body() {
13:         Block[][] stacks = table.getStacks();
14:         for(int i=0; i<stacks.length; i++) {
15:             for(int j=0; j<stacks[i].length; j++) {
16:                 Block block = (Block)getBeliefbase().getBeliefSet("blocks").getFact(stacks[i][j]);
17:                 Block target = stacks[i][j].getLower()==table
18:                 ? (Table)getBeliefbase().getBelief("table").getFact()
19:                 : (Block)getBeliefbase().getBeliefSet("blocks").getFact(stacks[i][j].getLower());
20:
21:                 IGoal stack = createGoal("stack");
22:                 stack.getParameter("block").setValue(block);
23:                 stack.getParameter("target").setValue(target);
24:                 dispatchSubgoalAndWait(stack);
25:             }
26:         }
27:     }
28: }

```

Figure 1.11. Java code for ConfigureBlocksPlan

“stack” plan is reused (see line 51) to move all blocks from the top of the block to be cleared to the table. To resolve the parameters used for body creation, a bindings declaration is used (lines 46-50). The variable \$upper is assigned to all blocks located on top of the given block (select statement in line 48). For each of these variable assignments an instance of the plan is created, assuring that all blocks are removed from the given block.

The two plan bodies are shown in Figs. 1.10 and 1.11, respectively. References to classes and methods provided by the Jadex engine are shown in boldface. Both plan classes define a constructor which takes the plan arguments and stores them in corresponding fields (lines 6-12 resp. 6-10) such that they are accessible from the body() methods, which will be described in turn.

The body() method of the StackBlocksPlan (Fig. 1.10, lines 14-24) first clears both blocks provided as arguments, and then moves the first block on top of the other. To clear the first block, a goal of type “clear” (cf. Fig. 1.9) is created (line 15) and the parameter is set to the block (line 16). The dispatchSubgoalAndWait() method (line 17) forces the

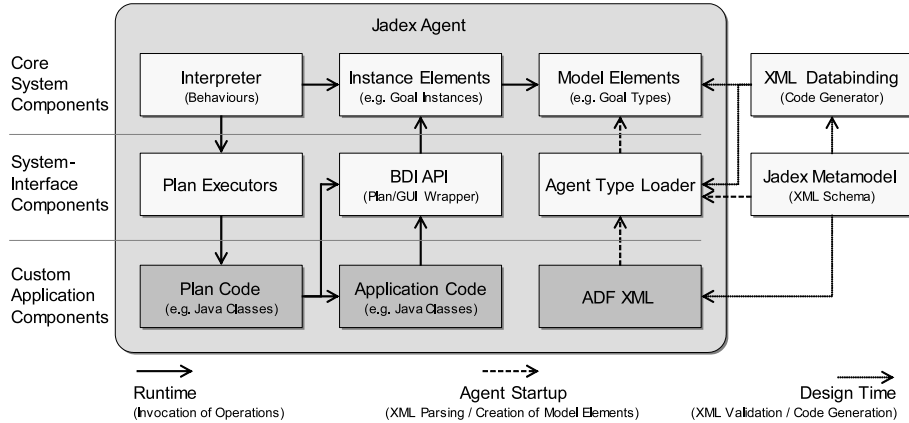


Figure 1.12. System realization

agent to adopt the goal, and halts the execution of the plan until goal processing is finished. If the goal fails, an exception is thrown leading to the whole plan to fail. Otherwise, the plan continues to clear the target block in a similar fashion (lines 19-21). Finally, the plan stacks the blocks on each other by calling the `stackOn()` method of the `Block` class (line 23).

In the `ConfigureBlocksPlan` (Fig. 1.11), the `body()` method (lines 12-25) consists of two loops through all stacks on the table, and all blocks of each stack, as returned by the `getStacks()` method of the `Table` class (line 13). This table object represents the desired target configuration. The agent now has to look up the corresponding blocks in its beliefbase, and then operate on this blocks such that they resemble the target configuration. The lookup is simple for the block itself, as the corresponding object can be directly from the belief set (line 16). The lookup of the object below the block (lines 17-19) is somewhat more difficult, because the block could be located directly on the table (line 18) or on top of another block (line 19). To perform the actual changes to the retrieved objects, a “stack” goal is created and dispatched (lines 21-24). Because the loop processes the stacks bottom-to-top, the sequential execution of all “stack” goals ensures that the final configuration resembles the desired target configuration.

1.4 Platform

This section describes the Jadex platform, i.e. the system realisation, and its integration into JADE. Figure 1.12 shows the essential com-

ponents required for developing and executing a Jadex agent, and highlights the dependencies between those components. The components are distinguished in *core system components* (upper row) which realize the reasoning engine, *system interface components* (middle row) that provide and define the access points to the system, and *custom application components* (lower row) which have to be supplied by the agent developer. The links between the components can be categorized in *runtime dependencies* (i.e. between components in the first two columns from the left), dependencies that only apply during the *agent startup* phase (see third column components), and dependencies resolved at *design time* (right column).

We will describe the components starting from the right. Jadex is based on a BDI *metamodel* defined in XML Schema (cf. Sect. ??). This schema is on the one hand used to validate the agent models specified in XML agent definition files (*ADF*). On the other hand, an *XML databinding* framework⁴ is used to generate Java classes for the elements of the metamodel and for reading model elements from XML. When an agent is instantiated the generated *agent type loader* reads the user supplied XML agent model and automatically creates the corresponding *model elements*.

From these model elements instances are continuously created at runtime, represented by *instance elements*. The main *interpreter* operates on the current instance elements and executes plans to handle events and goals. *Plan executors* are used to hide the details of plan implementation types from the system. As a default there is a plan executor for executing *plan code* written in Java. Plan code may access any other *application code* or third party libraries written in a suitable language. Both plan and application code have access to the reasoning engine through a *BDI API*. It is provided to plan and other application (e.g. GUI) code by wrappers that encapsulate the instance elements, and ensure proper synchronization and deadlock-avoidance when the API is called from the plans, or from external threads respectively.

For integration into JADE, the platform management tool (RMA) has been extended slightly to support launching of Jadex agents, by selecting the corresponding agent model with a file chosser. The Jadex interpreter itself is realized as a generic JADE agent. This generic agent loads the agent model supplied at startup, and initializes the reasoning engine according the settings given in the model (e.g. initial beliefs, goals, and plans). The functionalities corresponding to the execution

⁴JBind Java-XML Data Binding Framework, see <http://jbind.sourceforge.net/>

model components (cf. section 1.2.3), are implemented in separate JADE behaviours, automatically running inside the agent. Using a reference to the JADE agent object, Jadex plans have direct access to all operations of the JADE API as well (e.g. for handling of FIPA ACL messages).

1.4.1 Available tools and documentation

The system distribution contains complete documentation materials for quickstart and reference purposes. An introductory tutorial made up of several exercises shows the usage of basic system features in a step-by-step manner. Moreover, the distribution provides several example applications including their commented source code. A user guide provides a systematic overview of all features and also serves as a reference manual. In addition, Javadocs of the plan programming API and a reference to the metamodel defined in XML Schema are provided. The available tools are covered in a separate guide. Apart from the documentation material included in the distribution, there are publically available online tools kindly hosted by SourceForge.net, such as web forums for discussion and support requests, a database for bug-reports and feature requests, and a general mailing list with online archives.

As a Jadex agent is still a JADE agent, all runtime tools provided by the JADE platform such as Sniffer and Dummy agent can also be used with Jadex agents. To enable a comfortable testing of the internals of Jadex agents additional tool agents have been developed. In Fig. 1.13 an example application (marsworld) is depicted together with the logger and introspector tools in a typical debugging session. The BDI introspector (Fig.1.13 bottom left and right hand side) serves two purposes. First, it supports the visualization and modification of the internal BDI concepts thus allowing inspection and reconfiguration of an agent at runtime. Secondly, it simplifies debugging through a facility for the stepwise agent execution. In the step mode it is possible to observe and control each event processing and plan execution step having detailed control over the dispatcher and scheduler. Hence it can be easily figured out what plans are selected for an event or goal.

With the help of the logger (see Fig.1.13 on the top right) the agent's outputs can be directed to a single point of responsibility at runtime. In contrast to simple console outputs the logger agent preserves additional information about the output such as its time stamp and its source (the agent and method). Using these artefacts the logger agent offers facilities for filtering and sorting messages by various criteria allowing a personalized view to be created. In addition to the tools already included in the latest release, a tool for multi-agent application deployment is cur-

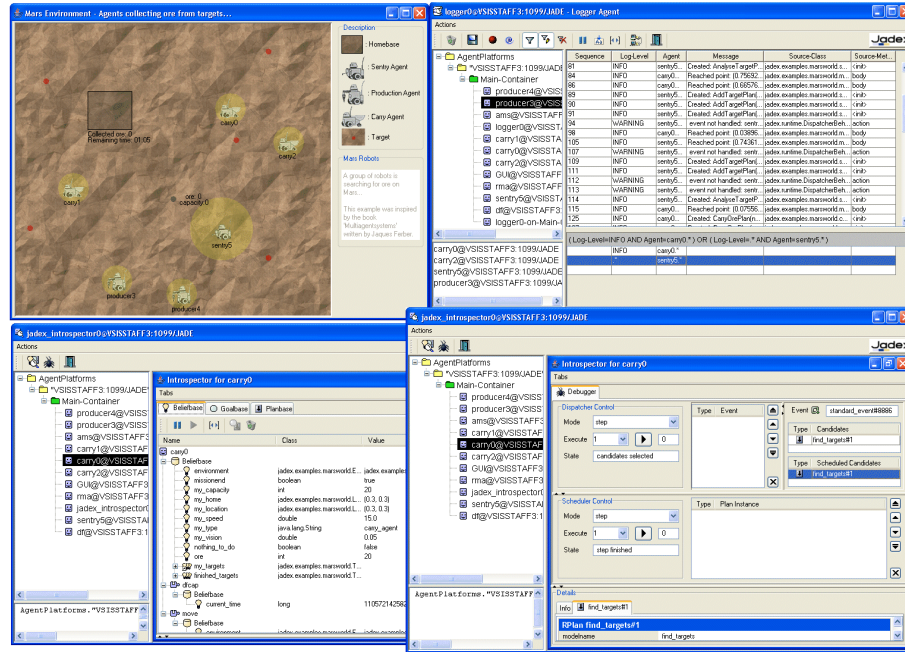


Figure 1.13. BDI introspector and logger screenshots

rently in development (see [4]), as well as a tool for on-line visualisation of agent execution based on ideas from [12].

1.4.2 Standards compliance, interoperability and portability

A driving factor for the development of Jadex was the need for a FIPA-compliant platform supporting advanced BDI reasoning capabilities. FIPA-compliance is achieved through the JADE platform, which provides sophisticated implementations of all important FIPA specifications. The Jadex reasoning engine, realized on top of the JADE platform, in itself only supports homogeneous (i.e. BDI) agents, but provides interoperability with agents based on other models. Agents realized using the conventional JADE programming techniques can be executed directly together with Jadex agents running on the same platform. Interoperability with other kinds of agents is straightforward as long as those agents adhere to the FIPA standard. E.g. in an example application, Jadex agents have been successfully connected to agents running on

the CAPA platform [10], which provides a petri-net based computation model for agents.

The reasoning engine is realized as a separate component, with only a few dependencies to the underlying platform, mainly the integration of the reasoning process to the platform execution model, and references to the handling of messages. The integration into the execution model involves, that methods of the message receiver, dispatcher and scheduler component (cf. Fig. 1.3) be called at appropriate times, what should be straightforward to implement in any agent programming environment. Therefore, although its current implementation is coupled to JADE, the reasoning engine can be easily integrated with other FIPA-compliant agent platforms such as CAPA [10] or ADK⁵, given that they provide a similar interface for message handling. It should also be possible in principle to use the system in conjunction with other middleware environments such as J2EE or .NET, when FIPA-compliance is not needed. As these environments usually do not provide support for active objects, a simple execution model would have to be developed (e.g. using a separate Java thread for each agent).

The engine was realized in Java 1.4 and includes the third party packages JBind for XML data binding and Apache Velocity⁶ for generating the content of some tool dialogs. To support mobile devices, a port of the engine is also available in a reduced version based on J2ME / CDC. Moreover, all kinds of tools and libraries with a Java API can easily be used to provide additional features. For example, in a larger project the Cayenne database mapping framework⁷ was used to connect agents to a relational database.

1.5 Applications supported by the language and/or the platform

Jadex is a general purpose development environment for creating multi-agent system applications, allowing to build agents with reactive (event-based) and deliberative (goal-driven) behaviour. It is not bound to a specific target domain, but has been used to realize applications in different domains such as simulation, scheduling, and mobile computation. Jadex originated in the MedPAge (“Medical Path Agents”) project [15, 16], which is part of the German priority research programme 1083 *Intelligent Agents in Real-World Business Applications* funded by the

⁵<http://www.tryllian.com/>

⁶<http://jakarta.apache.org/velocity/>

⁷<http://objectstyle.org/cayenne/>

Deutsche Forschungsgemeinschaft (DFG). In cooperation with the business management department of the University of Mannheim, the project investigates the advantages of using agent technology in the context of hospital logistics. In this project Jadex was used to realize a multi-agent application for market-based negotiation of treatment schedules [16], as well as for the simulation of a hospital model to test the negotiation mechanism [6]. In other contexts Jadex was used to realize portable PDA-based applications. A personal mobile task planner was developed, to test the Jadex J2ME port and to prove the usefulness of BDI agents on mobile devices [11]. Elsewhere, in the PITA (“Personal Intelligent Travel Assistant”) project at the Delft University of Technology, Jadex was used to realize a prototype of a mobile personal travel assistant application [2].

Besides building specific agent applications, Jadex has also been used for teaching and research regarding agent oriented software development in general. Due to its simple language based on well-known technologies such as Java and XML, and the extensive documentation material and simple example applications, Jadex is well-suited for teaching purposes. It has been successfully applied in several courses at the University of Hamburg, and is also evaluated by other institutes. Regarding research in agent systems, the project is also designed as a means for researchers to further investigate which mentalistic concepts are appropriate in the design and implementation of agent systems. The combination of XML Schema with Java databinding techniques allows the Jadex metamodel to be flexibly adapted and extended for experimentation purposes. While investigating different representations for beliefs, goals and plans, the system has been applied to several well-known AI problem domains (blocksworld, cleanerworld, mars robots, hunter-prey). These applications are also included in the distribution. Moreover, the Technical University of Karlsruhe has used Jadex to implement an experimental system for representing norms in multi-agent systems [21].

1.6 Final Remarks

In this chapter the Jadex BDI reasoning engine has been presented. The realization of the system is motivated mainly by three factors. Firstly, the system aims to combine the benefits of agent middleware and internal agent reasoning processes. Secondly, it intends to enhance the state-of-the-art BDI architecture by addressing some shortcomings of current BDI agent platforms such as implicit goal representation and thirdly, the system targets on making agent technology more easily us-

able by exploiting current software engineering techniques such as XML, Java and OQL.

The architecture of Jadex is in principle similar to traditional BDI systems, when event resp. goal processing is considered. Nevertheless, conceptual differences exist mainly concerning the representation of BDI core concepts and as well on language level. According to the usability requirement, beliefs are expressed in an object-oriented way instead of using logical formulae. Moreover, goals are represented as explicit durable entities instead of relying on events. On language level, Jadex differentiates between the description of runtime agent behaviour and its static structure. Therefore, for each of these purposes different languages are employed. The static agent structure is declared in an XML-dialect following the Jadex BDI metamodel specified in XML-schema, whereas ordinary Java is used for plan realization. BDI-specific facilities are made accessible from within plan through an application program interface.

Ongoing work currently focuses on two aspects of the system: Extensions to internal concepts and additional tool support. On the conceptual level extensions to the basic BDI-mechanisms are developed, such as support for planning, teams, and goal deliberation. It is planned to utilize the explicit representation of goals by improving the BDI architecture with a generic facility for goal deliberation which alleviates the necessity for designing agents with a consistent goal set. Additionally the explicit representation allows to investigate task delegation by considering goals at the inter-agent level. Work on tools mainly addresses the usability of agent technology as a mainstream software engineering paradigm. The tool support of Jadex currently focusses on the testing phase supplying a debugger and a logger agent. To achieve a higher degree of usability it is planned to support the design phase as well with a graphical modeling tool based on the MDA-approach [1]. Additionally, tools for documenting agents, visualizing agent behaviour within graphs and deployment of multi-agent applications are being developed [4].

The current version is Jadex 0.921, which can be freely downloaded under LGPL license from the project homepage <http://jadex.sourceforge.net/>. It is termed a beta stage release, what means that it has reached considerable stability and maturity to be used in experimental settings, but compatibility between releases is not guaranteed.

Acknowledgments

This work is partially funded by the German priority research programme 1083 *Intelligent Agents in Real-World Business Applications*.

Appendix: Summary

- 1(a) Reactive and deliberative behaviour is supported based on the BDI model and the corresponding mental attitudes. In addition to the basic BDI interpreter known from PRS systems, an explicit representation of goals is provided.
- 1(b) FIPA-compliant speech-act based communication is provided by the underlying JADE platform.
- 1(c) Weak mobility is provided by the underlying JADE platform. When developing mobile agents, some features of the system (e.g. thread-based plans) are not available.
- 1(d) The language is easy to learn, as it is based on well known technologies such as Java and XML. Experiments with students have shown that new users are quickly able to develop their first agents.
- 1(e) No formal semantics is available.
- 1(f) The language is very general and allows to create different kinds of agent applications.
- 1(g) The system does not define a new language for programming agent behaviour, but instead makes BDI-specific agent facilities available as application program interface (API). Hence, the BDI feature set can be easily extended.
- 1(h) No path to formal verification is provided.
- 1(i) The XML language enforces strong typing. The plan language inherits the software engineering and programming language principles of Java. In addition, reusability is supported by the definition of agent-modules called capabilities.
- 1(j).i Embedding the agent language into a general purpose language is not necessary, because the system cleanly separates the definition of an agent's structure and the definition of agent behaviour. The structure of an agent is defined in a system specific XML dialect following a BDI-metamodel, while the agent behavior is realized as plans coded directly in the general purpose programming language Java.
- 1(j).ii The default plan language is Java and therefore allows accessing any other application code or third party library written in Java. In addition, it is possible to define wrappers that allow to execute plans written in other (e.g. visual) languages.
- 2(a).i The documentation includes an introductory tutorial, a user guide which also serves as a reference manual, and a guide to the available tools. Javadocs of the plan programming API, and a reference to the metamodel defined in XML Schema are provided, and the distribution includes several example applications with source.
- 2(a).ii The system is based on JADE 3.2 and runs on any platform that supports Java 1.4. The distribution includes the third party packages JBind for XML databinding and Apache Velocity for generating the content of some tool dialogs.
- 2(b) The system complies with the FIPA-standards as implemented by JADE.
- 2(c) The system is Open Source and carefully designed and documented to allow easy and flexible extension of the provided functionality.

- 2(d).i** In addition to the tools provided by the JADE platform such as Sniffer and Dummy Agent, the system supplies tools to inspect the beliefs, goals, and plans of an agent at runtime, and to execute agents step-by-step. Moreover, a Logger Agent allows to collect, filter, and view logging output produced by the different agents of a multi-agent application.
- 2(d).ii** Apart from the documentation material included in the distribution (e.g. user guide and tutorial), there are publically available web forums for discussion and support requests, a database for bug-reports and feature requests, and a general mailing list with online archives.
- 2(d).iii** No additional tools (apart from those provided by JADE) are yet available, but a tool for multi-agent system deployment is currently in development.
- 2(e)** All kinds of tools and libraries with a Java API can be used within Jadex. For example, in a larger project the Cayenne database mapping framework was used to connect agents to a relational database.
- 2(f)** Although its current implementation is coupled to JADE, the reasoning engine is designed to be used on top of any existing middleware. It can be easily ported to other FIPA-compliant agent platforms such as CAPA or ADK and in principle to other middleware environments such as J2EE or .NET.
- 2(g).i** The performance of the system regarding the number of agents and messages is bounded by the performance of the underlying JADE platform. The computation cost induced by the reasoning engine highly depends on the complexity of the agents.
- 2(g).ii** The system is available as stable Open Source distribution and has already been used in several 3rd party projects. Nevertheless, the set of features is continuously evolving, and compatibility between releases is not guaranteed.
- 2(h).i** The system realizes a specific internal agent architecture, and therefore itself does not address heterogeneity, but it is possible to run Jadex agents on the same platform as any other JADE agents. Openness is supported in principle through FIPA-compliant communication, but not especially facilitated by the design of the system.
- 2(h).ii** Just like in JADE, agents are autonomous self-dependent entities with their own thread of control. A hierarchical structure of agents is not supported, but agents can be decomposed into hierarchically structured modules, which are similar to agents, but do not have their own reasoning process.
- 2(h).iii** The system includes a module for communication with a directory facilitator (DF) and for using simple FIPA interaction protocols (e.g. request).
- 3(a)** The system has been used mainly in research projects and teaching courses, e.g. to realize a multi-agent application for market-based negotiation of patient treatment dates, as well as for the simulation of a hospital model. The system has also been applied in mobile environments and to some well-known AI problem domains such as blocksworld and cleanerworld.
- 3(b)** The system is general purpose and not bound to a particular application domain.

References

- [1] H. Becker. Realisierung eines metamodellbasierten Entwurfswerkzeuges für BDI-Agentensysteme. Diplomarbeit, Distributed Systems and Information Systems Group, Computer Science Department, University of Hamburg/Germany, 2005. (in German).
- [2] M. Beelen. Personal Intelligent Travelling Assistant: a distributed approach. Master of science thesis, Knowledge Based Systems group, Delft University of Technology, 2004.
- [3] M. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, Massachusetts, 1987.
- [4] L. Braubach, A. Pokahr, K.-H. Krempels, and W. Lamersdorf. Deployment of Distributed Multi-Agent Systems. In *Fifth International Workshop on Engineering Societies in the Agents World (ESAW 2004)*, 2004.
- [5] L. Braubach, A. Pokahr, and W. Lamersdorf. Jadex: A Short Overview. In *Net.ObjectDays 2004: AgentExpo*, 2004.
- [6] L. Braubach, A. Pokahr, W. Lamersdorf, K.-H. Krempels, and P.-O. Woelk. A Generic Simulation Service for Distributed Multi-Agent Systems. In *From Agent Theory to Agent Implementation (AT2AI'04)*, 2004. to appear.
- [7] L. Braubach, A. Pokahr, D. Moldt, and W. Lamersdorf. Goal Representation for BDI Agent Systems. In *Proceedings of the Second Workshop on Programming Multiagent Systems: Languages, frameworks, techniques, and tools (ProMAS04)*, 2004.
- [8] P. Busetta, N. Howden, R. Rönquist, and A. Hodgson. Structuring BDI Agents in Functional Clusters. In N. R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI, Proceedings of the 6th International Workshop, Agent Theories, Architectures, and Languages (ATAL) '99*, pages 277–289. Springer, 2000.
- [9] D. Dennett. *The Intentional Stance*. Bradford Books, 1987.

- [10] M. Duvigneau, D. Moldt, and H. Rölke. Concurrent Architecture for a Multi-agent Platform. In F. Giunchiglia, J. Odell, and G. Weiß, editors, *Agent-Oriented Software Engineering III. Third International Workshop, AOSE 2002, Bologna, Italy, July 2002. Revised Papers and Invited Contributions*, volume 2585 of *LNCS*. Springer, 2003.
- [11] M. Harbeck. BDI-Agentensysteme auf mobilen Geräten. Diplomarbeit, Distributed Systems and Information Systems Group, Computer Science Department, University of Hamburg/Germany, 2004. (in German).
- [12] D. Lam and S. Barber. Debugging agent behavior in an implemented agent system. In *Second International Workshop on Programming Multi-Agent Systems at the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 45–56, 2004.
- [13] E. Mangina. Review of Software Products for Multi-Agent Systems. <http://www.agentlink.org/resources/software-report.html>, 2002.
- [14] E. Norling. Folk Psychology for Human Modelling: Extending the BDI Paradigm. In *Proceedings of in the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, Ney York, USA, July 2004.
- [15] T. O. Paulussen, N. R. Jennings, K. S. Decker, and A. Heinzl. Distributed Patient Scheduling in Hospitals. In G. Gottlob and T. Walsh, editors, *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*. Morgan Kaufmann, 2003.
- [16] T. O. Paulussen, A. Zöllner, A. Heinzl, A. Pokahr, L. Braubach, and W. Lamersdorf. Dynamic Patient Scheduling in Hospitals. In M. Bichler, C. Holtmann, S. Kirn, J. Müller, and C. Weinhardt, editors, *Coordination and Agent Technology in Value Networks*. GITO, Berlin, 2004.
- [17] A. Pokahr and L. Braubach. *Jadex User Guide, Release 0.921*, 2004.
- [18] A. Pokahr, L. Braubach, and W. Lamersdorf. Jadex: Implementing a BDI-Infrastructure for JADE Agents. *EXP – in search of innovation*, 3(3):76–85, 2003.
- [19] S. Poslad and P. Charlton. Standardizing Agent Interoperability: The FIPA Approach. In M. Luck et al., editor, *9th ECCAI Advanced Course, ACAI 2001 and Agent Links 3rd European Agent Systems Summer School, EASSS 2001, Prague, Czech Republic, July 2001*, pages 98–117. Springer-Verlag: Heidelberg, Germany, 2001.

- [20] A. Rao and M. Georgeff. BDI Agents: from theory to practice. In V. Lesser, editor, *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*, pages 312–319, San Francisco, CA, USA, 1995. The MIT Press: Cambridge, MA, USA.
- [21] T. Schubert. Normen zur Überwachung und Steuerung autonomer Multi-Agenten Systeme. Diplomarbeit, Institut für Programmstrukturen und Datenorganisation, Fakultät für Informatik, Universität Karlsruhe (TH), 2004. (in German).