

A Practical Agent Programming Language

Mehdi Dastani and John-Jules Ch. Meyer
Utrecht University
The Netherlands
{mehdi , j j}@cs . uu . nl

Abstract

This paper discusses the need for an effective and practical BDI-based agent-oriented programming language. It proposes an alternative by presenting the syntax and semantics of a programming language, called 2APL (A Practical Agent Programming Language). This programming language facilitates the implementation of multi-agent systems consisting of individual cognitive agents. 2APL distinguishes itself from other BDI-based agent-oriented programming languages by realising an effective integration of declarative and imperative style programming. This is done by introducing both declarative goals and events (which are used interchangeably in other programming languages) and by providing practical programming constructs.

1 Introduction

Existing BDI-based agent-oriented programming languages such as Jason[2], Jack[7], Jadex[6], and 3APL[4] provide programming constructs and mechanisms that allow direct implementation of software agents in terms of BDI concepts. These programming languages differ from each other as they facilitate the implementation of different but overlapping sets of agent concepts. For example, they all share programming constructs to support the implementation of an agent's beliefs and plans. However, 3APL differs from other programming languages as it supports the implementation of declarative goals as well as plan revision mechanism, but lacks programming constructs to support the implementation of events and event handling mechanism. Although, a comparison between these BDI-based programming languages is outside the scope of this paper, we would like to emphasize that some of the concepts that are shared by these languages have different semantics or even are not comparable at all. For example, the beliefs and goals in 3APL (and the beliefs in Jason) are propositions having declarative semantics while the beliefs in Jack and Jadex can be represented by conventional data structures lacking a declarative semantics. The situation gets even worse because different languages use different concepts with the same name or they use the same concept with different names. For example, although both Jason and 3APL use the concept of goal, a goal in Jason is an event that triggers a plan while a goal in 3APL is

a proposition that can be reasoned with [3, 5]. The declarative nature of goals in 3APL allows the implementation of generic planning rules that assign plans to *subgoals*.

Our experience with using these agent-oriented programming languages in various academic courses and research projects have shown that some of the existing programming constructs, tools, and mechanisms are indeed useful, expressive and effective in programming software agents. It also made it clear that new programming constructs, mechanisms and tools are needed to make the BDI-based programming languages even more expressive and *practical*. In our opinion, a challenge of a practical BDI-based agent-oriented programming language is an effective integration of declarative and imperative style programming. The declarative style programming should facilitate the implementation of the mental state of agents allowing agents to reason about their mental states and update them accordingly. The imperative style programming should facilitate the implementation of processes, the flow of control, and mechanisms such as procedure call, recursion, and interface to existing imperative programming languages.

In this paper, we present a BDI-based agent-oriented programming language called 2APL (A *Practical* Agent Programming Language). The main motivation to design and develop 2APL is an effective integration of programming constructs that support the implementation of declarative concepts such as belief and goals with imperative style programming such as events and plans. Like most BDI-based programming languages, different types of actions such as belief and goal update actions, test actions, external actions (to be performed in agents' environment), and communication actions are distinguished. These actions will be composed by conditional choice operators, iteration operator, and sequence operator. The composed actions constitute the plans of the agents. Like the existing agent programming languages, 2APL supports the so-called second principle planning by providing rules to indicate that certain goals can be achieved by certain pre-compiled plans. Agents may select and apply such rules to generate plans to achieve their goals.

As agents may operate in dynamic environments, they have to observe (be notified about) their environmental changes. In 2APL such environmental changes will be notified to the agents by means of *events*. It should be noted that in some agent programming languages events are used for various purposes. For example, in Jason events are used to model an agent's goals, while in Jadex events are generated to notify an agent's internal changes. In our view, although both goals and events cause an agent to execute actions, there are fundamental differences between them. For example, an agent's goal denotes a desirable state for which the agent performs actions to achieve it, while an event carries information about (environmental) changes which may cause an agent to react and execute certain actions. After the execution of actions, an agent's goal may be dropped if the state denoted by it is achieved, while an event can be dropped just before executing the actions that are triggered by it. Moreover, because of the declarative nature of goals, an agent can reason about its goals while an event only carries information which is not necessarily the subject of reasoning.

Some characterizing 2APL features are related to the constructs designed with respect to an agent's plans. The first construct is a part of an exception handling mechanism allowing a programmer to specify how an agent should repair its plans when their executions fail. This construct has the form of a rule which indicates how a plan should be repaired. It is similar to the plan revision rules introduced in 3APL, but it differs

from it as 2APL rules can only be applied to repair *failed* plans. In contrast, the plan revision rules of 3APL is applied to all plans continuously. In our view, it does not make sense to modify a plan if the plan is correct and executable. The second 2APL programming construct with respect to plans is related to the so-called critical (region of) plans. In most agent-oriented programming languages, an agent can have various plans whose executions can be interleaved. The arbitrary interleaving of plans may be problematic in some cases such that a programmer may want to indicate that a certain part of a plan should be executed at once without being interleaved with the actions of other plans.

2 2APL: Syntax

This section presents the complete syntax of 2APL, which is specified using the EBNF notation. In this specification, illustrated in Figure 1, we use $\langle atom \rangle$ to denote a Prolog like atomic formula starting with lowercase letter, $\langle Atom \rangle$ to denote a Prolog like atomic formula starting with a capital letter, $\langle ident \rangle$ to denote a string and $\langle Var \rangle$ to denote a string starting with a capital letter. We use $\langle ground_atom \rangle$ to denote a grounded atomic formula. Note that the question mark '?' used in $\langle testbelief \rangle$ is a syntactic entity of the 2APL language, while other question marks are used as elements of the EBNF notation indicating the optional choice.

An individual 2APL agent may be composed of various ingredients that specify different aspects of the agency. A 2APL agent can be programmed by implementing the initial state of those ingredients. The state of some of these ingredients will change during the agent's execution while the state of other ingredients remains the same during the whole execution of the agent. In the following, we will discuss each ingredients and give examples to illustrate them.

2.1 Beliefs and Goals

A 2APL agent may have beliefs and goals which change during the agent's execution. The *beliefs* of the agents are implemented by the belief base, which contains information the agent believes about its surrounding world including other agents. The implementation of the initial belief base starts with the keyword 'Beliefs:' followed by one or more belief expressions of the form $\langle belief \rangle$. Note that a $\langle belief \rangle$ expression is treated as a Prolog fact or rule such that the belief base of a 2APL agent becomes a Prolog program. All facts are assumed to be grounded.

```
Beliefs:  
pos(1,1).  
hasGold(0).  
trash(2,5).  
trash(6,8).  
clean(blockWorld) :- not trash(X,Y).
```

```

<2APL_Prog> ::= ("Include:" <ident>
| "Beliefupdates:" <BelUpSpec>
| "Beliefs:" <belief>
| "Goals:" <goals>
| "Plans:" <plans>
| "PG rules:" <pgrules>
| "PC rules:" <pcrules>
| "PR rules:" <prrules>)*
<BelUpSpec> ::= ("{" <query> "}" <beliefupdate> "{" < literals> "}")+
<belief> ::= ( <ground_atom> ":" | <atom> ":" -" < literals> ":" )+
<goals> ::= <goal> ("," <goal>)*
<goal> ::= <ground_atom> ("and" <ground_atom>)*
<baction> ::= "skip" | <beliefupdate> | <sendaction> | <javaaction> | <abstractaction>
| <testbelief> | <testgoal> | <adoptgoal> | <dropgoal>
<plans> ::= <plan> ("," <plan>)*
<plan> ::= <baction> | <sequenceplan> | <ifplan> | <whileplan> | <criticalplan>
<beliefupdate> ::= <Atom>
<sendaction> ::= "Send(" <iv> "," <iv> "," <atom> ")"
| "Send(" <iv> "," <iv> "," <iv> "," <iv> "," <atom> ")"
<javaaction> ::= "Java(" <iv> "," <atom> "," <Var> ")"
<abstractaction> ::= <atom>
<testbelief> ::= <query> "?"
<testgoal> ::= <query> "!"
<adoptgoal> ::= "adopta(" <goalvar> ")" | "adoptz(" <goalvar> ")"
<dropgoal> ::= "dropGoal(" <goalvar> ")"
| "dropSubgoal(" <goalvar> ")"
| "dropExactgoal(" <goalvar> ")"
<sequenceplan> ::= <plan> ";" <plan>
<ifplan> ::= "if" <query> "then" <scopeplan> ("else" <scopeplan>)?
<whileplan> ::= "while" <query> "do" <scopeplan>
<criticalplan> ::= "[" <plan> "]"
<scopeplan> ::= "{" <plan> "}"
<pgrules> ::= <pgrule>+
<pgrule> ::= <query>? "< -" <query> "|" <plan>
<pcrules> ::= <pcrule>+
<pcrule> ::= <atom> "< -" <query> "|" <plan>
<prrules> ::= <prrule>+
<prrule> ::= <planvar> "< -" <query> "|" <planvar>
<goalvar> ::= <literal> ("and" <literal>)*
<planvar> ::= <plan> | <Var> | "if" <query> "then" <scopeplanvar> ("else" <scopeplanvar>)?
| "while" <query> "do" <scopeplanvar> | <planvar> ";" <planvar>
<scopeplanvar> ::= "{" <planvar> "}"
< literals> ::= <literal> ("," <literal>)*
<literal> ::= <atom> | "not" <atom>
<ground_literal> ::= <ground_atom> | "not" <ground_atom>
<query> ::= "true" | <query> "and" <query> | <query> "or" <query> | "(" <query> ")" | <literal>
<iv> ::= <ident> | <Var>

```

The example above illustrates the implementation of the initial belief base of a 2APL agent. This belief base represents the information of an agent about its `blockworld` environment. In particular, the agent believes that its position in this environment is `(1,1)`, it has no gold in possession, there are trash at positions `(2,5)` and `(6,8)`, and that the `blockworld` environment is clean if there are no trash anymore.

The *goals* of a 2APL agent are implemented by its goal base, which consists of formulas each of which denotes a situation the agent wants to realize (not necessary all at once). The implementation of the initial goal base starts with the keyword `'Goals:'` followed by a list of goal expressions of the form $\langle goal \rangle$. Each goal expression is a conjunction of ground atoms. Note that a ground atom is treated as a Prolog fact. The following example is the implementation of the initial goal base of a 2APL agent. This goal base indicates that the agent wants to achieve a desirable situation in which it has five golds and the `blockworld` is clean. Note that this single conjunctive goal is different than having two separate goals `'hasGold(5)'` and `'clean(blockworld)'`. In the latter case, the agent wants to achieve two desirable situations independently of each other, e.g., first having a clean `blockworld` without any gold and then having five gold and perhaps a `blockworld` which is not clean anymore. Note that the separated goals in the goal base are separated by a comma.

```
Goals:
hasGold(5) and clean(blockworld)
```

The beliefs and goals of agents are related to each other. In fact, if an agent believes a certain fact, then the agent does not pursue that fact as a goal. This means that if an agent modifies its belief base, then its goal base may be modified as well.

2.2 Basic Actions

Basic actions specify the capabilities that an agent can perform to achieve its desirable situation. The basic actions will constitute an agent's plan, as we will see in the next subsection. In 2APL, six types of basic actions are distinguished: actions to update the belief base, communication actions, external actions to be performed in an agent's environment, abstract actions, actions to test the belief and goal bases, and actions to manage the dynamics of goals.

2.2.1 Belief Update Action

A *belief update action* updates the belief base of an agent when executed. A belief update action $\langle beliefupdate \rangle$ is an expression of the predicate argument form where the predicate starts with a capital letter. Such an action is specified in terms of pre- and post-conditions. An agent can execute a belief update action if the pre-condition of the action is derivable from its belief base. The pre-condition is a formula consisting of literals composed by disjunction and conjunction operators. The execution of a belief update action modifies the belief base in such a way that after the execution the post-condition of the action is derivable from the belief base. The post-condition of a belief update action is a list of literals. The update of the belief base by such an action

removes the atom of the negative literals from the belief base and adds the positive literals to the belief base. The specification of the belief update actions starts with the keyword 'Beliefupdates:' followed by the specifications of a set of belief update actions *<beliefupdatespecification>*.

```

Beliefupdates:
{not carry(gold)}      Pickup()      {carry(gold)}
{trash(X,Y) and pos(X,Y)} RemoveTrash() {not trash(X,Y)}
{hasGold(X)}          AddGold()      {not hasGold(X),
                                     not carry(gold),
                                     hasGold(X+1)}
{pos(X,Y)}            ChgPos(X1,Y1) {not pos(X,Y),
                                     pos(X1,Y1)}

```

Above is an example of the specification of the belief update actions. In this example, the specification of the Pickup() indicates that this belief update action can be performed if the agent does not already carry gold (i.e., the agent can carry only one gold item) and that after performing this action the agent will carry gold. Note that the agent cannot perform two Pickup() action consecutively. Note the use of variables in the specification of RemoveTrash(); it requires that an agent can remove trash if it is the same location as the trash. Note also that variables in the post-conditions are bounded since otherwise the facts in the belief base will not be grounded.

2.2.2 Communication Action

A *communication action* passes a message to another agent. A communication action *<sendaction>* can have either three or five parameters. In the first case, the communication action is the expression Send(Receiver, Performative, Language, Ontology, Content) where Receiver is a name referring to the receiving agent, Performative is a speech act name (e.g. inform, request, etc.), Language is the name of the language used to represent the content of the message, ontology is the name of the ontology used in the content of the message, and Content is an expression representing the content of the message. It is often the case that agents assume a certain language and ontology such that it is not necessary to pass them as parameters of their communication actions. The second version of the communication action is therefore the expression Send(Receiver, Performative, Content). It should be noted that 2APL interpreter is built on the Jade platform. For this reason, the name of the receiving agent can be a local name or a full Jade name. A full jade name has the form localname@host:port/JADE where localname is the name as used by 2APL, host is the name of the host running the agent's container and port is the port number where the agent's container, should listen to (see [1] for more information on Jade standards).

2.2.3 External Action

An *external action* is supposed to change the external environment in which the agents operate. The effects of external actions are assumed to be determined by the envi-

ronment and might not be known to the agents beforehand. The agent thus decides to perform an external action and the external environment determines the effect of the action. The agent can come to know the effects of an external action by performing a sense action, defined as an external action, or by means of events generated by the environment. An external action $\langle javaaction \rangle$ is an expression of the form `Java(Env, ActionName, Return)`. The parameter `Env` is the name of the agent's environment, implemented as a Java class. The parameter `ActionName` is a method call (of the Java class) that specifies the effect of the external action in the environment. The environment is assumed to have a state represented by the instance variables of the class. The execution of an action in an environment is then a read/write operation on the state of the environment. The parameter `Return` is a list of values, possibly an empty list, returned by the corresponding method. An example of the implementation of an external action is `Java(blockworld, east(), .)` (go one step east in the blockworld environment). The effect of this action is that the position of the agent in the blockworld environment is shifted one slot to the right. No return value is expected.

2.2.4 Abstract Action

An *abstract action* is an abstraction mechanism allowing the encapsulation of a plan by a single action. An abstract action will be instantiated with a concrete plan when the action is executed. The instantiation of plans with abstract actions are specified through special rules called PC-rule, which stands for procedure call rules (see subsection 2.4.2 for a description of PC-rules). In fact, the general idea of an abstract action is similar to a procedure call in imperative programming languages while the PC-rules function as procedure definitions. Like a procedure call, an abstract action $\langle abstractaction \rangle$ is an expression of the predicate argument form starting with a lowercase letter.

2.2.5 Belief and Goal Test Actions

A *belief test action* is to test whether a belief expression is derivable from an agent's belief base, i.e., it tests whether the agent has a certain belief. A belief test action $\langle testbelief \rangle$ is an expression of the predicate argument form followed by a question mark. Such a test may generate a substitution for the variables that are used as arguments in the belief expression. A belief test action is basically a (Prolog) query to the belief base which can be used in a plan 1) to instantiate a variable in rest of the plan, or 2) to block the execution of the plan (if the test fails). A *goal test action* is to test whether a formula is derivable from the goal base, i.e., whether the agent has a certain goal. Like a belief test action, this action can be used to instantiate a variable with a value, or to block the execution of the rest of a plan. A goal test action $\langle testgoal \rangle$ is an expression of the predicate argument form followed by an exclamation mark.

2.2.6 Goal Dynamics Actions

The *adopt goal* and *drop goal* actions are used to adopt and drop a goal to and from the agent's goal base, respectively. The adopt goal action $\langle adoptgoal \rangle$ can have two different forms: `adopta(ϕ)` and `adoptz(ϕ)`. These two actions can be used to add

the goal ϕ (a conjunction of literals) to the begin and to the end of an agent's goal base, respectively. Note that the programmer has to ensure that the variables in ϕ are instantiated before these actions are executed since the goal base should always be grounded. Finally, the drop goal action $\langle dropgoal \rangle$ can have three different forms: $dropGoal(\phi)$, $dropSubGoal(\phi)$, and $dropExactGoal(\phi)$. These actions can be used to drop from an agent's goal base, respectively, all goals that are a subgoal of ϕ , all goals that have ϕ as a subgoal, and exactly the goal ϕ .

2.3 Plans

In order to reach its goals, a 2APL agent adopts *plans*. A plan consists of basic actions composed by process operators. In particular, basic actions can be composed by means of the sequence operator, conditional choice operators, conditional iteration operator, and an unary operator to identify critical (region of) plans.

The sequence operator $;$ is a binary operator generating the plan $\langle sequenceplan \rangle$ from two other plans. It indicates that the first plan should be performed before the second plan. The conditional choice operator generates the plan $\langle ifplan \rangle$, which is an expression of the form $if \phi \text{ then } \pi_1 \text{ else } \pi_2$. The condition of such an expression (i.e., ϕ) is evaluated with respect to an agent's belief base. This expression can thus be interpreted as to perform the if-part of the plan (i.e., π_1) when the agent believes ϕ , otherwise the agent performs the else-part of the plan (i.e., π_2). The conditional iteration operator generates the plan $\langle whileplan \rangle$ which is an expression of the form $while \phi \text{ do } \pi$. The condition ϕ should also be evaluated with respect to an agent's belief base. This iteration expression is then interpreted as to perform the plan π in the body of the while loop as long as the agent believes ϕ . The last unary operator generates the plan $\langle criticalplan \rangle$ which is an expression of the form $[\pi]$. This plan is interpreted as a critical plan π , which should be executed at once ensuring that the execution of π is not interleaved with the actions of other plans. Note that an agent can have different plans at the same time.

The plans of a 2APL agent are implemented by its plan base. The implementation of the initial plan base starts with the keyword 'Plans:' followed by a list of plans. The following example illustrates the implementation of the initial plan base of a 2APL agent. The first plan is a critical plan ensuring that the agent updates its belief base with its new position $(X+1, Y)$ immediately after performing a belief test action followed by an external `east` action in the `blockworld` environment. This guarantees that no external action from other plans of the agent (e.g., the external `south` action from the second plan in the plan base) can be interleaved, i.e., it guarantees that the update of the agent's position after the `east` action in the critical plan is a correct update.

Plans:

```
[pos(X, Y)?; Java(blockWorld, east(), -); ChgPos(X+1, Y)],
  pos(X, Y)?; Java(blockWorld, south(), -); ChgPos(X, Y+1)
```

2.4 Reasoning Rules

The 2APL programming language provides constructs to implement practical reasoning rules that can be used to implement the generation of plans. In particular, three types of practical reasoning rules are proposed: planning goal rules, procedural call rules, and plan repair rules. In the following subsections, we explain these three types of rules.

2.4.1 Planning Goal Rules (PG rules)

A planning goal rule can be used to implement that an agent should generate a plan if it has certain goals and beliefs. The specification of a planning goal rule $\langle pgrule \rangle$ consists of three entries: the head of the rule, the condition of the rule, and the body of the rule. The head and the condition of a planning goal rule are query expressions used to check if the agent has a certain goal and belief, respectively. The body of the rule consists of a plan in which variables may occur. These variables should be bounded by the variables that occur in the goal and belief expressions. A planning goal rule of an agent can be applied when the goal and belief expressions (in the head and the condition of the rule) are derivable from the agent's goal and the belief bases, respectively. The application of a planning goal rule involves an instantiation of variables that occur in the head and condition of the rule as they are queried from the goal and belief bases. The resulted substitution will be applied to the generated plan to instantiate its variables. A planning goal rule is of the form:

$\langle query \rangle ? \text{''} < - \text{''} \langle query \rangle \text{''} | \text{''} \langle plan \rangle$

Note that the head of the rule is optional which means that the agent can generate a plan only based on its belief condition. The following is an example of a planning goal rule indicating that a plan to go to a position $(X2, Y2)$ departing from a position $(X1, Y1)$ in order to remove trash at $(X2, Y2)$ can be generated if the agent has the goal `clean(blockworld)` and it believes there is trash at position $(X2, Y2)$. Note that `goTo(X1, Y1, X2, Y2)` is an abstract action (see subsection 2.4.2 for how to execute an abstract action).

PG-rules:

```
clean(blockworld) <- pos(X1,Y1) and trash(X2,Y2) |
{goTo(X1, Y1, X2, Y2);RemoveTrash()}
```

Note that this rule can be applied if (beside the satisfaction of the belief condition) the agent has a conjunctive goal `hadGold(5)` and `clean(blockworld)` since the head of the rule is derivable from this goal.

2.4.2 Procedural Call Rules (PC rules)

The procedural call rules is introduced for various reasons and purposes. In fact, such a rule is introduced to generate plans as a response to the reception of messages send by other agents, events generated by the external environment, and the execution of abstract actions. Like planning goal rules, the specification of procedural call rules consist of three entries. The only difference is that the head of the procedural rules is an atom

$\langle atom \rangle$ (predicate-argument expression), rather than a query $\langle query \rangle$, which represents a message, an event, or an abstract action. A message and an event are represented by atoms with the special predicates `message` and `event`, respectively. An abstract action is represented by any predicate name starting with a lowercase letter. Note that like planning goal rules, a procedural call rule has a belief condition indicating when a message (or event or abstract action) should cause the generation of a plan. Thus, a procedural call rule can be applied if the agent has received a message (or an event or executes an abstract action) and the belief query of the rule is derivable from its belief base. The instantiation of variables and application of the resulting substitutions to the plan variables are the same as with planning goal rules. A procedural call rule $\langle pcrule \rangle$ is of the form:

$\langle atom \rangle \text{ " < - " } \langle query \rangle \text{ " | " } \langle plan \rangle$

The following are examples of procedural call rules. The first rule indicates that if an agent A informs that there is some gold at position (X2, Y2) and the agent believes it does not carry any gold, then the agent has to go from its current position to the gold position, pick up the gold, go to the depot position (i.e. position (3, 3)), and drop the gold in the depot. The `PickUp()` and `AddGold()` are belief update actions to administrate the facts that the agent is carrying gold and has certain amount of gold, respectively. The second rule indicates that if the environment notifies the agent that there is some gold at position (X2, Y2) and the agent believes it does not carry gold, then the abovementioned sequence of actions should take place. The generation of plans without a belief condition enables a programmer to implement reactive agent behaviour, i.e., plans are generated if the agent is notified about an environmental change. Finally, the last rule indicates that the abstract action `goTo` should be performed as a certain sequence of actions. Note that all plans are implemented as critical plans. The reason is that in these plans external and belief update actions will be executed consecutively such that an unfortunate interleaving of actions can have undesirable effects. Note the use of recursion in this PC-rule.

PC-rules:

```
message(A, inform, La, On, goldAt(X2, Y2))
  <- not carry(gold) |
  {[ pos(X1, Y1)?;
     goTo(X1, Y1, X2, Y2);
     Java(blockworld, pickup(), _);
     PickUp();
     goTo(X2, Y2, 3, 3);
     Java(blockworld, drop(), _);
     AddGold()
  ]}

event(gold(X2, Y2), blockworld, me)
  <- not carry(gold) |
  {[ pos(X1, Y1)?;goTo(X1, Y1, X2, Y2);
     Java(blockworld, pichup(), _);PickUp();
     goTo(X2, Y2, 3, 3);
```

```

        Java(blockworld, drop(), _) ; AddGold()
    ]}

goTo(X1, Y1, X2, Y2) <- X1 < X2 |
    {[ Java(blockworld, east(), _) ; ChgPos(X1+1, Y1) ;
        goTo(X1+1, Y1, X2, Y2)
    ]}

```

2.4.3 Plan Repair Rules (PR rules)

Like other practical reasoning rules, a plan repair rule consists of three entries: two abstract plan expressions and one belief query expression. We have used the term abstract plan expression since such plan expressions include variables that can be instantiated with a plan. A plan repair rule indicates that if the execution of an agent's plan (i.e., any plan that can be instantiated with the abstract plan expression) fails and the agent has a certain belief, then the failed plan should be replaced by another plan. A plan repair rule $\langle prrule \rangle$ has the following form:

$\langle planvar \rangle \text{ " < - " } \langle query \rangle \text{ " | " } \langle planvar \rangle$

A plan repair rule of an agent can thus be applied if 1) the execution of one of its plan fails, 2) the failed plan can be matched with the abstract plan expression in the head of the rule, and 3) the belief query expression is derivable from the agent's belief base. The satisfaction of these three conditions results in a substitution for the variables that occur in the abstract plan expression in the body of the rule. Note that some of these variables will be instantiated with a part of the failed plan through the match between the abstract plan expression in the head of the rule and the failed plan. For example, if π, π_1, π_2 are plans and X is a plan variable, then the abstract plan $\pi_1; X; \pi_2$ can be matched with the failed plan $\pi_1; \pi; \pi_2$ resulting the substitution $X = \pi$.

The resulted substitutions will be applied to the second abstract plan expression to generate the new (repaired) plan. The following is an example of a plan repair rule. This rule indicates that if the execution of a plan that starts with the external action $\text{Java}(\text{blockworld}, \text{east}(), L)$ fails, then the plan should be replaced by a plan in which the agent first goes one step to north, then makes two steps to east, and goes one step back to south. This repair can be done without a specific belief condition. Note the use of the variable X that indicates that any failed plan starting with $\text{Java}(\text{blockworld}, \text{east}(), L)$ can be repaired by the same plan in which the first external action is replaced by four external actions.

```

PR-rules: Java(blockworld, east(), _) ; X <- true |
    {Java(blockworld, north(), _) ;
      Java(blockworld, east(), _) ;
      Java(blockworld, east(), _) ;
      Java(blockworld, south(), _) ; X}

```

The question is when the execution of a plan fails. We consider the execution of a plan as failed if the execution of its first action fails. When the execution of an action fails depends on the type of action. The execution of a belief update action fails if the

action is not specified, an abstract action if there is no applicable procedural call rule, an external (Java) action if the environment throws an `ExternalActionFailedException`, an belief test action if the belief expression is not derivable from the belief base, a test goal action if the goal expressions are not derivable from the goal base, and a critical plan section if one of its actions fails. The execution of all other actions will always be successful. When the execution of an action fails, then the execution of the whole plan is stopped. The failed action will not be removed from the failed plan such that it can be repaired.

2.5 External Environment

An agent can perform actions in different external environments. The environments are shared between all agents. Any Java class that implements the *Environment* interface can be used as a 2APL environment. The *Environment* interface contains two methods, *addAgent(String name)* to add an agent to the environment and *removeAgent(String name)* to remove an agent from the environment. The constructor of the environment must require exactly one parameter of the type *ExternalEventListener*. This object listens to external events.

The execution of action *Java(env, f(a₁, . . . , a_n), R)* calls a method with name *f* in environment *env* with $n \geq 1$ arguments. The first argument of this function is the identifier of the agent that executes the action. The environment needs to have this identifier, for example, to pass information back to the agent by means of events. The last parameter *R* of the external action is meant to pass information back to the *plan* in which the external action was executed. Note that the execution of a plan is halted until the method *f* is ready and the return value is accessible to the rest of the plan. If the return type of the method is *Term*, *R* will be bounded to the result of the method. *R* remains unbounded otherwise.

Methods may throw an *ExternalActionFailedException*. If they throw an exception, the corresponding external action is considered as failed. The following is an example of a method that can be called as external action.

```
public Term move(String agent, String direction)
    throws ExternalActionFailedException
{
    if (direction.equals("north") {moveNorth();}
    else if (direction.equals("east") {moveEast();}
    else if (direction.equals("south") {moveSouth();}
    else if (direction.equals("west") {moveWest();}
    else throw
        new ExternalActionFailedException("Unknown direction");

    return getPositionTerm();
}
```

2.6 Events and Exception

In 2APL information can also be passed through *events* and *exceptions*. The main use of events is to pass information from environments and agents to agents. When implementing an environment in Java, the programmer should decide when and which information from the environment should be passed to agents. This can be done in an environment implemented in Java by calling the method `notifyEvent(AF event, String... agents)` in the *ExternalEventListener* which was an argument of the environments constructor. The first argument of this method may be any valid atomic formula. The rest of the arguments may be filled with strings that represents local names of agents. The events can be caught by agents whose name is listed in the argument list to trigger one of their procedural call rules. If the programmer does not specify any agents in the argument list, all agents can catch the event. Such a mechanism of generating events by the environment and catching it by agents can be used to implement the agents' perceptual mechanism. Moreover, sending messages and executing abstract actions also generate and throw events that subsequently are caught by the agents to trigger their procedural rules. In practice, the generation and throwing of events for these two actions will be performed by the 2APL interpreter and not by 2APL programmers.

The use of exceptions in 2APL is limited to the application of plan repair rules. In fact, a plan repair rule is triggered when a plan execution fails. Exceptions are used to notify that the execution of a plan was not successful. The exception contains the identifier of the failed plan such that it can be determined which plan needs to be repaired. 2APL does not provide programming constructs to implement the generation and throwing of exceptions. Like events for the send and abstract actions, exceptions are semantical entities that cannot be used by 2APL programmers.

3 2APL: Semantics

In the previous section, we described 2APL programming constructs and their intuitive meanings. In this section, we present the formal semantics of 2APL in terms of a transition system. A transition system is a set of derivation rules for deriving transitions. A transition is a transformation of one configuration into another and it corresponds to a single computation step. Because of the space limitation, we only present the configuration of 2APL agents, external actions, and characterizing 2APL constructs such as goal related constructs, critical plan construct, and plan repair rules.

The configuration of an individual agent consists of its identifier, beliefs, goals, plans, specifications of belief update actions, reasoning rules, the substitutions resulted from queries to the belief and goal bases, and the received events. Since reasoning rules and the specification of belief update actions, do not change during an agent's execution, we will not include them in the agent's configuration to keep the presentation as simple as possible. It should be noted that additional information is assigned to an agent's plan. In particular, an identifier is assigned to each plan which can be used to notify that the execution of the plan is failed. This is needed to identify and repair the plans the execution of which have failed. Moreover, the instantiation of the PG rule

through which a plan is generated is assigned to the plan. This information is used to avoid selecting a PG rule to generate a plan if there is still a plan in the plan base generated based on the same goal condition (not belief condition). This means that the application of planning rules are indifferent to belief changes.

definition 1 *The configuration of an individual 2APL agent is defined as $A_i = \langle \iota, \sigma, \gamma, \Pi, \theta, \xi \rangle$ where ι is a string representing the agent's identifier, σ is a set of belief expressions (belief) representing the agent's belief base, γ is a list of goal expressions (goal) representing the agent's goal base, Π is a set of plan entries representing the agent's plan base, θ is a ground substitution that binds domain variables to domain terms, and ξ is the agent's event base. Each plan entry is a tuple (π, r, p) where π is the executing plan, r is the instantiation of the PG rule through which π is generated, and p is the plan identifier. The agent's event base ξ is a set of tuples $\langle E, I, M \rangle$ where*

- *E is the set of events thrown by external environments. Such an event has the form event(A, S, R) where A is an atomic formula carrying the information passed by an environment, S is the identifier of the environment, and R is the identifier of the desired receiver for this event (ι in this case).*
- *I is the set of exceptions thrown because of a plan execution failure. The exception is the identifier of the failed plan in the plan base.*
- *M is the set of messages sent to the agent. A message is of the form message(s, r, p, l, o, ϕ) where s is the identifier of the sender of the message, r is the identifier of the receiver of the message, p is the performative of the message, l is the language of the message, o is the ontology of the message, and ϕ is an atom representing the message content.*

In the rest of this paper, we use \models as a first-order entailment relation (we use Prolog engine for the implementation of this relation). We also assume $\sigma \models \phi \Leftrightarrow \gamma \not\models \phi$, i.e., an agent's belief can never be the agent's goal since a goal entailed by the belief base implies that the goal is achieved and thus is not a goal anymore.

The configuration of a multi-agents system is defined in terms of the configuration of individual agents in the multi-agent system and their shared external environments.

definition 2 *Let A_i be the configuration of agent i and let χ be a set of external shared environments each of which is a set of atoms (atom). The configuration of a 2APL multi-agents system is defined as $\langle A_1, \dots, A_n, \chi \rangle$*

The execution of an external action $Java(env, \alpha(t_1, \dots, t_n), V)$ has two different effects. The shared environments will be changed and the variable V might become bounded to a Term. To define the effect of an external action on the agent's state we define a function that returns a tuple containing the new environment and the binding for V . Let $F_\alpha^{env}(t_1, \dots, t_n, \chi)$ be the function that executes external action α with arguments t_1, \dots, t_n in the environment $env \in \chi$ and returns a tuple (τ, χ') , where τ contains one substitution for the output variable V and χ' is the updated set of environments (a change in env may change other environments in χ). A successful execution of an external action updates the agent's substitution Θ and the set of shared environments

χ . Note that because the environment is shared among agents, the transition for an external action of an individual agent is defined at the multi-agent level.

$$\frac{F_{\alpha}^{env}(t_1\theta, \dots, t_n\theta, \chi) = (t, \chi')}{\langle A_1, \dots, A_i, \dots, A_n, \chi \rangle \longrightarrow \langle A_1, \dots, A'_i, \dots, A_n, \chi' \rangle}$$

where

$$A_i = \langle i, \sigma_i, \gamma_i, \{(\text{Java}(\text{env}, \alpha(\mathbf{t}_1, \dots, \mathbf{t}_n), \mathbf{V}), \mathbf{r}, \text{id})\}, \theta, \xi \rangle \&$$

$$A'_i = \langle i, \sigma_i, \gamma_i, \{(\epsilon, \mathbf{r}, \text{id})\}, \theta \cup \{\mathbf{V}/t\}, \xi \rangle$$

However, if the execution of an external action fails, then F_{α}^{env} returns (\perp, χ') and the environment generates an exception. Note that the failed action remains in the plan base, the environments χ may be updated, and the event base ξ is updated to capture the failure exception.

$$\frac{F_{\alpha}^{env}(t_1\theta, \dots, t_n\theta, \chi) = (\perp, \chi')}{\langle A_1, \dots, A_i, \dots, A_n, \chi \rangle \longrightarrow \langle A_1, \dots, A'_i, \dots, A_n, \chi' \rangle}$$

where

$$A_i = \langle i, \sigma_i, \gamma_i, \{(\text{Java}(\text{env}, \alpha(\mathbf{t}_1, \dots, \mathbf{t}_n), \mathbf{V}), \mathbf{r}, \text{id})\}, \theta, \langle E, I, M \rangle \rangle \&$$

$$A'_i = \langle i, \sigma_i, \gamma_i, \{(\text{Java}(\text{env}, \alpha(\mathbf{t}_1, \dots, \mathbf{t}_n), \mathbf{V}), \mathbf{r}, \text{id})\}, \theta, \langle E, I \cup \{\text{id}\}, M \rangle \rangle$$

In order to achieve an agent's goal, plans should be generated and executed. The generation of plans is through application of planning goals rules. A PG rule $\kappa \leftarrow \beta \mid \pi$ specifies that plan π can achieve goal κ . Applying PG rules will only update the plan base. Let P be the set of all possible plans, I be the set of all plan identifiers, and $\kappa \leftarrow \beta \mid \pi$ be a variant of any PG rule, i.e., all variables occurring in the rule are assumed to be fresh variables.

$$\frac{\begin{array}{l} \gamma = [\gamma_1, \dots, \gamma_i, \dots, \gamma_n] \& \gamma_i \models \kappa\tau_1 \& \sigma \models \beta\tau_1\tau_2 \& \\ \neg\exists\pi', \pi'' \in P, \neg\exists\text{id} \in I : (\pi', (\kappa\tau_1 \leftarrow \beta\tau_1\tau_2 \mid \pi''), \text{id}) \in \Pi \end{array}}{\langle \iota, \sigma, \gamma, \Pi, \theta, \xi \rangle \longrightarrow \langle \iota, \sigma, \gamma, \Pi', \theta, \xi \rangle}$$

Where $\Pi' = \Pi \cup \{(\pi\tau_1\tau_2, \kappa\tau_1 \leftarrow \beta\tau_1\tau_2 \mid \pi\tau_1\tau_2, \text{id})\}$. Note that it is checked that there is not already a plan in Π which is generated by the same planning rule for the same goal.

Goals can also be adopted and dropped from the goal base by means of specific $\text{adoptX}(\mathbf{g})$, $\text{adoptz}(\mathbf{g})$ and $\text{dropGoal}(\mathbf{g})$ actions. The first two actions add goal \mathbf{g} to the goal base.

$$\frac{\sigma \not\models \mathbf{g}\theta}{\langle \iota, \sigma, \gamma, \{(\text{adoptX}(\mathbf{g}), \mathbf{r}, \text{id})\}, \theta, \xi \rangle \longrightarrow \langle \iota, \sigma, \gamma', \{(\epsilon, \mathbf{r}, \text{id})\}, \theta, \xi \rangle}$$

where $\gamma' = [\mathbf{g}\theta \mid \gamma]$ if adoptX is adopta (indicating that the goal $\mathbf{g}\theta$ is added to the begin of the list γ of goals) and $\gamma' = [\gamma \mid \mathbf{g}\theta]$ if adoptX is adoptz (indicating that the goal $\mathbf{g}\theta$ is added to the end of γ). The $\text{dropGoal}(\mathbf{g})$ action drops all goals that are subgoals of \mathbf{g} from the goal base.

$$\frac{\gamma' = \gamma - \{f \mid \mathbf{g} \models f\}}{\langle \iota, \sigma, \gamma, \{(\text{dropGoal}(\mathbf{g}), \mathbf{r}, \text{id})\}, \theta, \xi \rangle \longrightarrow \langle \iota, \sigma, \gamma', \{(\epsilon, \mathbf{r}, \text{id})\}, \theta, \xi \rangle}$$

The transitions for $\text{dropSubGoal}(\mathbf{g})$ and $\text{dropExactGoal}(\mathbf{g})$ are similar except that $\gamma' = \gamma - \{f \mid f \models \mathbf{g}\}$ and $\gamma' = \gamma - \{f \mid f \equiv \mathbf{g}\}$, respectively. See section 2.2.6 for their intuitive meanings.

The execution of a critical plan is the execution of the maximum number of actions that occur in the plan in one step. Suppose the first failed action in the critical plan $[\alpha_1, \dots, \alpha_n]$ is α_m for $m \leq n$. An agent can go from its configuration A_1 to configuration A_{m+1} by applying one transition rule. Thus, we must check that all the transitions for α_1 to α_{m-1} are derivable and the transition for α_m is first one which is not derivable. To do this we define the following predicate:

$transition(A_i, A_{i+1}) =$
 $\langle \iota, \sigma_i, \gamma_i, \{([\alpha_i; \dots; \alpha_n], r, id)\}, \theta_i, \xi_i \rangle \longrightarrow \langle \iota, \sigma_{i+1}, \gamma_{i+1}, \{([\alpha_{i+1}; \dots; \alpha_n], r, id)\}, \theta_{i+1}, \xi_{i+1} \rangle$
 If all initial actions in the critical plan can be executed till α_m , then $transition(A_i, A_{i+1})$ must be derivable for all $1 \leq i \leq m$, but not for $i = m + 1$.

$$\frac{1 \leq m \leq n \ \& \ (\forall_i : 1 \leq i \leq m \rightarrow transition(A_i, A_{i+1})) \ \& \ \neg transition(A_{m+1}, A)}{\langle \iota, \sigma_1, \gamma_1, \{([\alpha_1; \dots; \alpha_n], r, id)\}, \theta_1, \xi_1 \rangle \longrightarrow \langle \iota, \sigma_{m+1}, \gamma_{m+1}, \{([\alpha_{m+1}; \dots; \alpha_n], r, id)\}, \theta_{m+1}, \xi_{m+1} \rangle}$$

Finally, the execution of the application of a plan repair rule is based on the received exceptions that identify failed plans. Let $\xi = \langle E, I, M \rangle$ be the event base of a 2APL agent and $\pi_1 \leftarrow \beta \mid \pi_2$ be a variant of a PR rule. Suppose the execution of a plan $(\pi, r, id) \in \Pi$ fails such that $id \in I$. Then, the plan repair rule can be applied if the failed plan π matches the abstract plan expression π_1 in the head of the rule, and moreover, its belief condition is derivable from the belief base. The result is a substitution that will be applied to the abstract plan expression in the body of the rule to generate a new plan and to add it to the plan base. We assume a unification function $U(\pi, \pi_1)$ that implements a prefix matching strategy for matching plan π with abstract plan expression π_1 . Roughly speaking, a prefix matching strategy means that the abstract plan expression is matched with the prefix of the failed plan. The unification function returns a tuple (τ_T, τ_P, π^*) where τ_T is a term substitution, τ_P is a plan substitution and π^* is the postfix of π that did not play a role in the match with π_1 (e.g., $U(\alpha(a); \alpha(b); \alpha(c), X; \alpha(Y)) = ([Y/b], [X/\alpha(a)], \alpha(c))$). Note the all substitutions are applied to the abstract plan expression from the body of the rule to generate a new plan.

$$\frac{U(\pi, \pi_1) = (\tau_T, \tau_P, \pi^*) \ \& \ \sigma \models \beta \tau_T \tau_P \ \& \ id \in I}{\langle \iota, \sigma, \gamma, \{(\pi, r, id)\}, \theta, \langle E, I, M \rangle \rangle \longrightarrow \langle \iota, \sigma, \gamma, \{(\pi_2 \tau_T \tau_P; \pi^*, r, id)\}, \theta, \langle E, I \setminus \{id\}, M \rangle \rangle}$$

If the head of no plan repair rule can be matched with the failed plan (i.e., unification is undefined for all rules, \perp), then the exception is deleted from the event base and the failed plan remains in the plan base.

$$\frac{\forall (\pi_1 \leftarrow \beta \mid \pi_2) \in PR : (U(\pi, \pi_1) = \perp \text{ or } \sigma \not\models \beta) \ \& \ id \in I \ \& \ (\pi, r, id) \in \Pi}{\langle \iota, \sigma, \gamma, \Pi, \theta, \langle E, I, M \rangle \rangle \longrightarrow \langle \iota, \sigma, \gamma, \Pi, \theta, \langle E, I \setminus \{id\}, M \rangle \rangle}$$

4 Conclusion and Future Works

In this paper, we presented a BDI-based agent-oriented programming language that provides practical constructs for the implementation of cognitive agents. The complete syntax and the intuitive interpretation of the involved programming constructs are discussed. Unfortunately, because of the space limitation we could only present the transition semantics of some characterising programming constructs.

We have implemented this semantics in the form of an interpreter that can execute 2APL programs (i.e., initial configuration of 2APL agents). The execution of agents is based on a deliberation cycle. Each cycle determines which transition in which order should take place. The 2APL interpreter starts with applying a planning goal rule to generate a plan for one of the agent's goal, selects and executes plans, checks for exceptions and repairs failed plans by applying plan repair rules, and finally checks for received messages and events to apply the procedural call rules. This interpreter is integrated in a 2APL platform that allows an agent programmer to load, edit, run, and debug a set of 2APL agents. This platform is built on top of the Jade platform in order to exploit all tools and facilities that are developed for the Jade platform. These include tools such as the Sniffer, Introspector, and RMA (Remote Agent management). We use also the Jade communication layer to implement the communication between agents. Note that the Jade platform aims to be compliant with the FIPA standards. Since the communication between 2APL agents are through the Jade platform, the 2APL interpreter inherits the objective of the Jade platform of being FIPA compliant.

We are working on various extensions of both 2APL language (e.g., adding constructs to implement organisations and coordination artifacts at the multi-agent level) as well as tools to be integrated in the 2APL platform (e.g., visual programming and debugging facilities).

References

- [1] F. Bellifemine, F. Bergenti, G. Caire, and A. Poggi. JADE - a java agent development framework. In *Multi-Agent Programming: Languages, Platforms and Applications*. Kluwer, 2005.
- [2] R. Bordini, J.F.Hübner, and R. Vieira. Jason and the golden fleece of agent-oriented programming. In *Multi-Agent Programming: Languages, Platforms and Applications*. Kluwer, 2005.
- [3] R. H. Bordini, M. Dastani, J. Dix, and A. El Fallah Seghrouchni. *Multi-Agent Programming: Languages, Platforms and Applications*. Springer, Berlin, 2005.
- [4] M. Dastani, M. van Riemsdijk, and J.-J. C. Meyer. Programming multi-agent systems in 3apl. In *Multi-Agent Programming: Languages, Platforms and Applications*. Kluwer, 2005.
- [5] M. Dastani, M. B. van Riemsdijk, and J.-J. C. Meyer. Goal types in agent programming. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI'06)*, 2006.
- [6] A. Pokahr, L. Braubach, and W. Lamersdorf. Jadex: A BDI reasoning engine. In *Multi-Agent Programming: Languages, Platforms and Applications*. Kluwer, 2005.
- [7] M. Winikoff. JACKTM intelligent agents: An industrial strength platform. In *Multi-Agent Programming: Languages, Platforms and Applications*. Kluwer, 2005.