Model checking agent programming languages

**Louise A. Dennis · Michael Fisher ·
Matthew P. Webster · Rafael H. Bordini**

Abstract In this paper we describe a verification system for multi-agent programs. This is the first *comprehensive* approach to the verification of programs developed using programming languages based on the BDI (belief-desire-intention) model of agency. In particular, we have developed a specific layer of abstraction, sitting between the underlying verification system and the agent programming language, that maps the semantics of agent programs into the relevant model-checking framework. Crucially, this abstraction layer is both flexible and extensible; not only can a variety of different agent programming languages be implemented and verified, but even *heterogeneous* multi-agent programs can be captured semantically. In addition to describing this layer, and the semantic mapping inherent within it, we describe how the underlying model-checker is driven and how agent properties are checked. We also present several examples showing how the system can be used. As this is the first system of its kind, it is relatively slow, so we also indicate further work that needs to be tackled to improve performance.

L.A. Dennis (✉) · M. Fisher
Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK
e-mail: l.a.dennis@liverpool.ac.uk

M. Fisher
e-mail: mfisher@liverpool.ac.uk

M.P. Webster
Virtual Engineering Centre, Daresbury Laboratory, Warrington WA4 4AD, UK
e-mail: matt@liverpool.ac.uk

R.H. Bordini
Institute of Informatics, Federal University of Rio Grande do Sul, PO Box 15064,
91501-970 Porto Alegre, RS, Brazil
e-mail: R.Bordini@inf.ufrgs.br

1 Introduction

Since the introduction of the *agent-based systems* concept in the 1980s (Bond and Gasser 1988; Bratman et al. 1988; Davis and Smith 1983; Cohen and Levesque 1990; Durfee et al. 1989; Shoham 1993) the field has seen significant growth and increasing research maturity. This is true not only within academia but also for industrial applications, where the agent metaphor has been shown to be useful in capturing many practical situations, particularly those involving complex systems comprising flexible, autonomous, and distributed components (Klügl et al. 2005).

The large number of agent platforms now available (Bordini et al. 2005a) has meant that the *industrial* uptake of this technology is continually growing. In software development, we have seen significant commercialisation of multi-agent systems technology, e.g., in the form of a Java-based ontology development and multi-agent toolkit (Himoff et al. 2005). Similarly, building effective and user-friendly transportation systems is increasingly tackled using AI methodologies and multi-agent technology in particular (Klügl et al. 2005). Other application areas have seen the emergence of agent frameworks, financed by industry, and designed to cope specifically with industrial requirements (Hirsch et al. 2008), and such technology has been successfully introduced in many companies (e.g., goal-oriented, autonomic process navigation; Rimassa and Burmeister 2007). An early survey of applications of agent technology in industrial systems control can be found in Jennings and Wooldridge (1998). Other areas in which agents are actively used include air-traffic control (Ljunberg and Lucas 1992), autonomous spacecraft control (Muscettola et al. 1998), and health care (Moreno and Garbay 2003). Clearly, many of these are areas for which we must demand software *dependability* and *security*.

As agent-based solutions are used in increasingly complex and critical areas, there is clearly a need to analyse, comprehensively, the behaviour of such systems. Not surprisingly, therefore, the area of formal verification techniques tailored specifically for agent-based systems is now attracting a great deal of attention (Fisher et al. 2007). While program verification is well advanced, for example Java verification using Java PathFinder (Visser et al. 2003; Java PathFinder 2009), the verification of agent-oriented programs poses new challenges that have not yet been adequately addressed, particularly within the context of practical model-checking tools. Tackling this deficiency is our long-term aim in this work.

We here consider the formal verification, via model checking, of programs written in Agent Programming Languages (APLs). As will be described below, such languages are significantly different from traditional programming languages and, hence, their verification requires extended techniques. The predominant model for agent systems is the BDI (belief-desire-intention) model (Rao and Georgeff 1992, 1995) within which agents are viewed as being *rational* and acting in accordance with their beliefs and goals. As such it is important to be able to verify properties expressed in terms of such concepts. Thus, in agent verification, we have to verify not only what the agent does, but *why* it chose that course of action, *what* it believed that made it choose to act in this way, and *what* its intentions were that led to this. Our previous work in this area has concentrated on one particular language, AgentSpeak (Rao 1996), whereas we here extend to *any* agent language that is based on the

BDI paradigm (Rao and Georgeff 1995) and for which a formal operational semantics can be provided within our framework. In addition, we now go beyond previous work in verifying more complex multi-agent systems. Though the performance is still not high when a larger number of agents are considered, we show that our approach is very flexible and, with the speed improvements envisaged, can potentially provide an effective verification system for much larger multi-agent systems.

In essence, our framework consists of two components. The *Agent Infrastructure Layer* (AIL) is a set of Java classes designed to act as a toolkit for creating interpreters for BDI Agent Programming languages. This toolkit is designed to make the construction of such interpreters quick and easy once an operational semantics is provided. The second component is *Agent JPF* (AJPF), a version of the Java Pathfinder (JPF) model checker (Visser et al. 2003; Java PathFinder 2009) which has been extended with a property specification language appropriate for agent programs and some Java interfaces suitable for encapsulating multi-agent systems in an efficient fashion.

Our aim has been to define a sufficiently general intermediate layer on which model checking could be performed. AIL and AJPF are designed to work efficiently together: AIL is designed to optimise model-checking time rather than execution time and to implement the interfaces specified in AJPF.[1] We are thus able to model check the properties of any (potentially heterogeneous) multi-agent system implemented in languages that have AIL-based interpreters. By optimising the AIL's state representation for use with AJPF, our customised version of the JPF model checker, we avoid the intricacies of carrying out such optimisation for each language separately, thereby reducing the probability of introducing errors at this stage.

The work described in this paper builds on results from a number of workshop (Dennis et al. 2008a; Webster et al. 2009) and conference (Bordini et al. 2008) papers. It develops and updates these results and includes unpublished applications (Webster et al. 2009) developed to assess the practicality of the software. This paper constitutes the first complete description of the final system, which is now available on Sourceforge, at http://mcapl.sourceforge.net.

2 Background

2.1 Rational agents and multi-agent systems

The agent abstraction captures truly *autonomous* behaviour. As we do not consider arbitrary, 'random' autonomy, then we assert that an autonomous system should have some *motivation* for acting in the way it does. This aspect is captured by the concept of a *rational agent* (Bratman 1987; Cohen and Levesque 1990; Rao and Georgeff 1992). The key aspect of a rational agent is that the decisions it makes, based on dynamic motivations, should be both "reasonable" and "justifiable". In detail, a BDI agent comprises *beliefs* that the agent has about itself and its environment, *desires* (or *goals*) representing its long-term aims, and *intentions* describing the

---

[1] Note that, if required, AIL can be used without AJPF and *vice versa*.

agent's immediate goals (the ones it is currently trying to achieve through acting on the environment where it is situated). Thus, such an agent analyses data about itself and its environment and generates, or updates, its beliefs. The agent endeavours to tackle its long-term aims within this context, which leads to a set of immediate goals (*intentions*), which are tackled through the agent's plans (determining appropriate courses of action to achieve goals). When several possible (conflicting) goals exist, the agent must undertake some deliberation and decide which intention to realise in practise as well as reasoning about how to do so.

A multi-agent system (MAS) is a system consisting of a number of rational agents interacting with each other. Problem solving using multi-agent systems is now an established area of software engineering. The cooperation aspect helps solve problems that are difficult to solve by individual agents or traditional, centralised, computer systems. Areas in which multi-agent systems have been successful include online trading, space missions, air traffic control, disaster response, and the modelling of social structures (Klügl et al. 2005; Hirsch et al. 2008; Rimassa and Burmeister 2007; Moreno and Garbay 2003), to name just a few.

2.2 BDI agent programming languages

There are *many* different agent programming languages and agent platforms based, at least in part, on the BDI approach. Particular languages developed for programming *rational* agents include AgentSpeak (Rao 1996), *Jason* (Bordini et al. 2005b, 2007), 3APL (Hindriks et al. 1998; Dastani et al. 2005), DRIBBLE (van Riemsdijk et al. 2003), Jadex (Pokahr et al. 2005), GOAL (Hindriks et al. 2001; de Boer et al. 2007), CAN (Winikoff et al. 2002), SAAPL (Winikoff 2007), GWENDOLEN (Dennis and Farwer 2008), and METATEM (Fisher and Hepple 2009; Fisher and Ghidini 2010). Rather than providing a framework in which the complex logical properties of systems using just *one* particular agent approach can be verified, we have developed a flexible and uniform framework allowing the verification of a wide range of agent-based programs, produced using several different high-level agent programming languages. Specifically, our focus has been on BDI languages, i.e., languages that generally follow the *beliefs*, *desires*, and *intentions* paradigm (Rao and Georgeff 1995) of agent-oriented programming. Consequently, the architecture presented in this paper is based upon our study of common concepts and structures appearing in the operational semantics of various BDI programming languages (Dennis et al. 2008a). Agents programmed in these languages commonly contain a set of *beliefs*, a set of *goals*, and a set of *plans*. Plans determine how an agent acts based on its beliefs and goals; the sophisticated form or plan execution is the basis for *practical reasoning* (i.e., reasoning about actions) in such agents. As a result of executing a plan, the beliefs and goals of an agent may change as the agent performs actions in its environment.

In Dennis et al. (2008a), we tackled issues in the treatment of beliefs, goals, plans as well as events, intentions, and other components central to the design of these languages. We started by extracting the major concepts of 3APL (Dastani et al. 2005) and the variant of AgentSpeak (Rao 1996) encapsulated by *Jason* (Bordini et al. 2007) for use in our framework. However, our approach does not exclude other languages,

in some cases even those based on completely different agent architectures. It was our aim to include languages that have practical relevance. Thus, we did not want to restrict ourselves to (abstract) programming languages that could not be considered for serious software development projects.

2.3 Agent verification

The main approaches to detailed systems analysis can be broadly categorised as *testing* (Hierons et al. 2009), *model checking* (Clarke et al. 1999), and *theorem proving* (Owre and Shankar 2008).

Testing is fast, but generally non-exhaustive, and thus cannot guarantee that a property holds throughout a given system. However, the latter two techniques are exhaustive and will, in principle, give a definite answer to the question of whether a property holds for all program executions. Theorem proving requires a deep knowledge of mathematical structures and techniques, is usually only partly automatable, and is, in practise, often very costly, since the task can usually only be carried out by experts. Model checking, on the other hand, relies on a complete, but automated, inspection of the system's state space. This makes it computationally costly (in terms of time and space), but does not usually require a specialist as in the case of theorem proving. In order to reduce the computational problems, state-of-the-art model checkers employ a number of reduction techniques that make it possible to handle even large state spaces relatively efficiently. Since we require an automatic/definite answer and are looking into practical applications of verification in the design of multi-agent programs, model checking is the most promising of the approaches.

Model checking is a technique whereby a finite description of a system is analysed with respect to a property in order to ascertain whether *all* possible executions of the system satisfy this property. Formally, the property is typically described using a temporal logic formula (Emerson 1990), while the model checking process essentially involves trying to find a model within the system description for the *negation* of this property. If any such model is found then it describes a 'run' of the system that does *not* satisfy the required property (Clarke et al. 1999). In agent-based systems, it is vital to verify not only the behaviour that the system has, but also to verify *why* the agents are undertaking certain courses of action within the multi-agent system. Thus, in our approach the temporal basis of model-checking, capturing the *dynamic* nature of agent computation, is extended with *modal operators* (Gabbay et al. 2003) capturing the *informational* ('beliefs'), *motivational* ('desires') and *deliberative* ('intentions') aspects of rational agents.

With the increasing use of multi-agent systems and the sophistication of model-checking technology, it is not surprising that the concept of automated agent verification has attracted significant interest. Pioneering work on *model checking* techniques for the verification of agent-based systems has appeared, for example, in Bordini et al. (2004), Kacprzak et al. (2004), Bordini et al. (2006), Raimondi and Lomuscio (2007). Our previous work (Bordini et al. 2004, 2006) has concentrated on model checking techniques for agent-based systems written in the logic-based agent-programming language AgentSpeak (Rao 1996). This required a specific encoding of the AgentSpeak agent system's states in the input language used by the model checker. Applying

similar techniques to other agent programming languages would require manually encoding the state representation of that language. This is both tedious and error-prone. This was a key motivation for the work presented here. Our intention was to lift away the effort involved in developing a model checking framework for a given language from encoding the language directly in a model checker, moving up to the level of constructing a (model-checker backed) interpreter for the language given suitable support.

3 AIL: a framework for creating BDI interpreters

The *Agent Infrastructure Layer* (AIL) is an intermediate layer that encompasses key concepts from a wide range of BDI programming languages as data structures in Java and enables the implementation of their operational semantics within a clear framework. These data structures can then be used both in creating Java interpreters for the agent programming languages and for interfacing with the underlying agent model checker.

An agent originally programmed in some agent programming language (APL) and running in an AIL-based interpreter uses the AIL data structures to store its internal state comprising, typically: a belief base, a plan library, a set of intentions, and other temporary state information. We also assume that the APL defines a *reasoning cycle* which is expressed using an operational semantics. The rules in the cycle define how the agent's practical reasoning progresses, depending on its current internal representation and the current stage of the reasoning cycle. The AIL provides support for constructing reasoning cycles along with a number of rules that commonly appear in the operational semantics of agent programming languages. The operational semantic rules within AIL are given in Appendix A. In short, the AIL toolkit collects together Java classes that:

1. facilitate the implementation of interpreters for various agent programming languages;
2. contain adaptable, clear, operational semantics; and
3. can be verified through AJPF, an extended version of the open source Java model checker JPF (Visser et al. 2003).

If an AIL-based interpreter is run in conjunction with the AJPF model checker, then the system will notify the model checker each time a new state is reached that is relevant to the verification. It is left to the designers of interpreters to decide where, in the reasoning cycle, such points should fall.

As a natural consequence of this, AIL also makes it easier to develop an interpreter for a programming language using the AIL classes than it would be to build an interpreter "from scratch" in Java. Figure 1 provides a diagrammatic representation of AIL within the AJPF model checking architecture.

AJPF is our extension of the JPF model checker (Visser et al. 2003; Java PathFinder 2009), which includes interfaces linking AIL interpreters to the model checking framework and a property specification language. These interfaces also allow programming languages that do *not* have their own AIL-based interpreters to
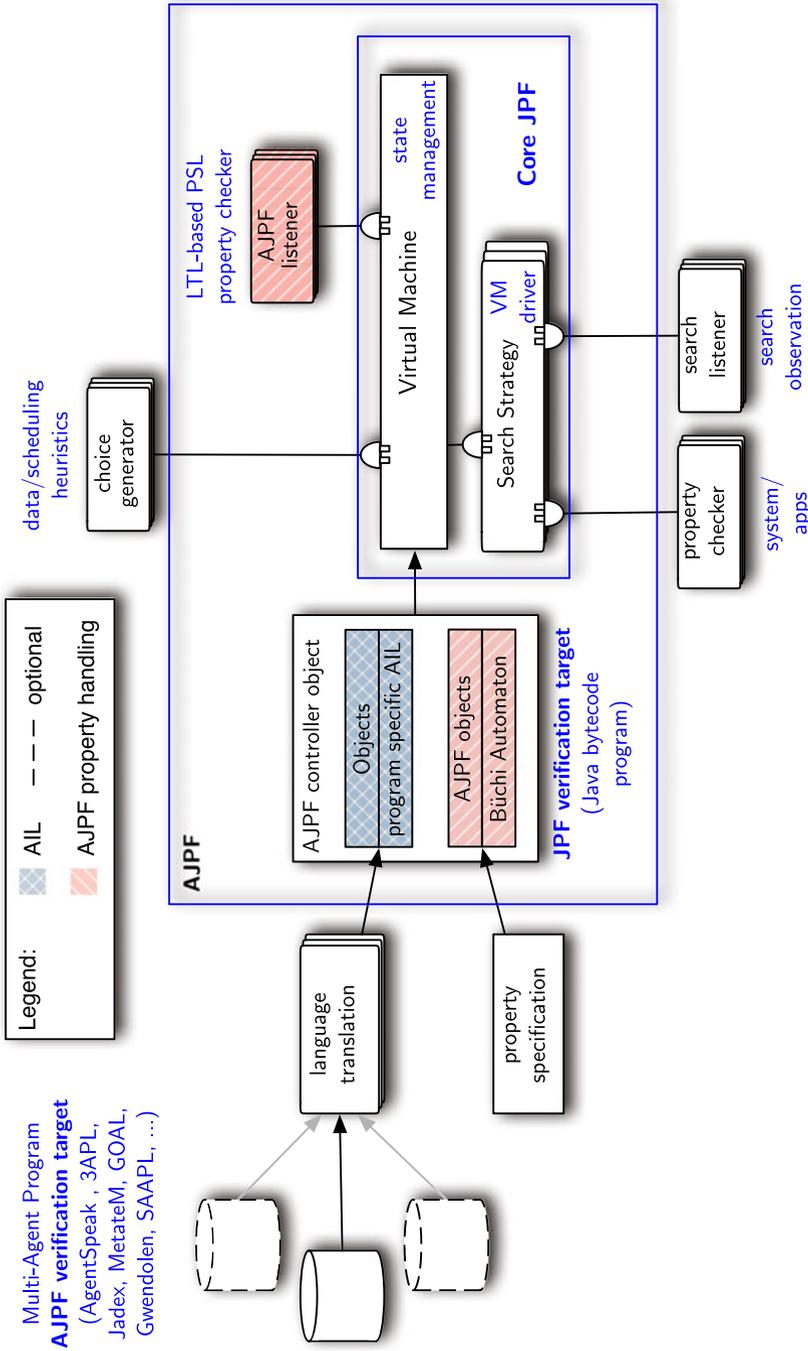
Fig. 1 AJPF architecture (Bordini et al. 2008)

be model checked against specifications written in the same property specification language, using AJPF. However these languages will not benefit from the efficiency improvements that the optimised AIL classes can provide. AJPF is discussed in more detail in Sect. 4.2.

Figure 1 shows how the combination of the translated agent program(s) together with the property specification constitutes an AJPF controller object. This controller is used as the JPF verification target. The original program(s) with the original property specification can be viewed as the AJPF verification target; they are fed into the appropriate translators available as part of our framework.

3.1 Components of agent programming languages

We will now proceed to give an overview of the AIL structures available.

3.1.1 Agents

Agents are specified by a large class with multiple fields. This class is intended for sub-classing or wrapping by language-specific interpreters which need only refer to the fields that are of interest to them. We discuss some of the key components of this class below. It also contains components for exploring organisational and group structures as outlined in Dennis et al. (2008b), though we do not discuss these here.

When we have discussed many of the key data structures within an agent, we will provide a formal definition of an AIL agent in Sect. 3.1.12.

3.1.2 Beliefs

Every agent contains a *belief base* of statements it believes to be true. This can be viewed as a set of first order formulae. AIL also supports Prolog-style reasoning based on a *rule base* of Horn clauses.

3.1.3 Guards and logical consequence

AIL uses the concept of a *guard* on a plan (as well as on other constructs). In the case of a plan, the guard is used to determine the applicability of the plan. In general, guards are used whenever an agent needs to check its own state.

A guard consists of two sets of expressions. These sets represent propositions the agent must hold (positive states)—i.e., statements it believes, goals it possesses etc.—and propositions the agent must not hold (negative states).[2] A guard represents the conjunction of these expressions which can be checked for validity against the agent's internal state (so typically such expressions refer to things the agents believes or goals it possesses). AIL provides an interface for the expressions that may appear in guards and requires these expressions to implement a *logical consequence* relation (we will write $ag \models g$ to represent that the guard, g, is a logical consequence of

---

[2]The system distinguishes between strong and weak negation, i.e. the difference between not believing something is the case and believing something is not the case.

the internal state of agent, *ag*). Procedurally speaking, the implementation of logical consequence provides a decision procedure for deciding whether a guard holds in an agent's current state.

AIL provides a default implementation of logical consequence which provides an algorithm for deciding whether some statement is believed by the agent, is a goal of the agent, has been sent as a message by the agent, as well as common propositional logical statements constructed from such atomic formulae (e.g. conjunction, disjunction, and negation).

However, since different languages may have different semantics for belief, goal, etc., this procedure can be overridden so, for instance, a different semantics can be supplied for what it means for a formula to be a goal of the agent. It is also possible, by this means, for individual languages to provide custom data structures for additional parts of an agent's state that those languages need to appear in guards.

The default implementation of logical consequence supports reasoning using Prolog-style Horn clauses where the literals appearing in such clauses may, themselves, be guards. This allows such rules to be used in languages with varying semantics for the literals that appear in such rules, without each language having to implement it's own version of Prolog-style reasoning.

The use of the logical consequence method for reasoning about an agent's internal state is a key aspect of the AIL system and it provides the bridge into the agent state for the atomic properties of the AJPF property specification language. Then, if such properties refer, for instance, to the fact that some agent believes something, the semantics for this is determined by the relevant agent's implementation of the logical consequence relation.

3.1.4 Goals

Goals are represented by a first order literal with a particular type—in the AIL toolkit the types are *achieve*, *perform*, *test* and *maintain* (following Dastani et al. 2005). There are rules in the toolkit distinguishing between the first three types—the fourth is currently unsupported. There is no reason why further goal types should not be added.

Agents do not maintain an explicit list of goals. Instead they deduce their goals from the "commit to goal" events that are found in their set of intentions. However, the 'Agent' class provides a number of methods that allow the agent to be treated as if it *did* maintain an explicit set of goals. In reality these methods inspect the intention data structure.

3.1.5 Intentions

AIL's most complex data structure is that which represents an *intention*. BDI languages use intentions to store the *intended means* for achieving goals—this is generally represented as some from of *deed stack* (deeds include actions, belief updates, and the commitment to goals). Intention structures in BDI languages may also maintain information about the (sub-)goal they are intended to achieve or the event that triggered them. In AIL, we aggregate this information: an intention becomes a stack

of tuples of an event, a guard, a deed, and a unifier. This AIL intention data structure is most simply viewed as a matrix structure consisting of four columns in which we record events (new perceptions, goals committed to and so forth), deeds (a plan of future actions, belief updates, goal commitments, etc.), guards (which must be true before a deed can be performed) and unifiers. These columns form an event stack, a deed stack, a guard stack, and a unifier stack. Rows associate a particular deed with the event that has caused the deed to be placed on the intention, a guard which must be believed before the deed can be executed, and a unifier. New events are associated with an empty deed, ε.

Example The following shows the full structure for a single intention to clean a room. We use a standard BDI syntax: !g to indicate the goal g, and +!g to indicate the commitment to achieve that goal (i.e., a new goal that g becomes true is adopted). Constants are shown starting with lower case letters, and variables with upper case letters.

Event	Deed	Guard	Unifier
`+!clean()`	`+!goto(Room)`	`dirty(Room)`	`Room = room1`
`+!clean()`	`+!vacuum(Room)`	⊤	`Room = room1`

This intention has been triggered by a desire to clean—the commitment to the goal *clean*() is the trigger event for both rows in the intention. An intention is processed from top to bottom so we see here that the agent first intends to commit to the goal *goto*(*Room*), where *Room* is to be unified with *room*1. It will only commit to this goal if it believes the (guard) statement, *dirty*(*Room*). Once it has committed to that goal it then commits to the goal *vacuum*(*Room*). In many languages the process of committing to a goal causes an expansion of the intention stack, pushing more deeds on it to be processed. So *goto*(*Room*) may be expanded *before* the agent commits to vacuuming the room. In which case the above intention might become

Event	Deed	Guard	Unifier
`+!goto(Room)`	`+!planRoute(Room, Route)`	⊤	`Room = room1`
`+!goto(Room)`	`+!follow(Route)`	⊤	`Room = room1`
`+!goto(Room)`	`+!enter(Room)`	⊤	`Room = room1`
`+!clean()`	`+!vacuum(Room)`	⊤	`Room = room1`

At any moment, we assume there is a *current intention* which is the one being processed at that time. The function $\mathscr{S}_{\text{int}}$ (implemented as a method in AIL) may be used to select an intention. By default, this chooses the first intention from a queue, but this choice may be overridden for specific languages and applications. Intentions can be *suspended* which allows further heuristic control. A suspended intention is, by default, *not* selected by $\mathscr{S}_{\text{int}}$. Typically an intention will remain suspended until some trigger condition occurs, such as a message being received. Many operational semantic rules (such as those involved with perception) resume *all* intentions—this allows suspension conditions to be re-checked.

3.1.6 Events

Events are things that occur within the system to which an agent may wish to react. Typically we think of these as changes in beliefs or the new commitment to goals. In many (though not all) programming languages, events trigger plans (i.e., a plan might selected for execution only when the corresponding event has taken place).

In AIL there is a special event, 'start', that is used to start off an intention which is not triggered by anything specific. This is mainly used for the initial goals of an agent—the intention begins as a start intention with the deed to commit to a goal. In some languages the belief changes caused by perception are also treated in this way. Rather than being added directly to the belief base, in AIL such beliefs are assigned to intentions with the event start and then added to the belief base when the intention is actually executed.

3.1.7 Deeds

Deeds appear in the bodies of plans and as stacks in intentions, and represent things the agent is planning to do but has not yet actually done. Deeds include:

Beliefs: an agent may both plan to add or to drop a belief.

Goals: an agent may both plan to commit to a goal (a sub-goal) or to drop a goal.

Actions: an agent may plan to perform an action.

No Plan Yet: the "no plan yet" deed is used when an intention contains an event that has yet to be planned for; we represent the "no plan yet" deed with the distinguished symbol, ε.

Lock/Unlock: a deed can lock or unlock an intention; the idea here is to allow an operational semantics to, under some conditions, force an intention to remain current until it is unlocked.

Waiting: the *wait for* deed allows intentions to be suspended until a particular guard is satisfied in the agent's state. We represent "waiting for guard g", as '$*g$'. The idea is to allow an operational semantics to remove an intention from consideration by $\mathscr{S}_{\text{int}}$ until some condition is met.

3.1.8 Plan library

The Agent class also contains a plan library. Plans are matched against intentions and/or the agent's state and manipulate existing (or create new) intentions. There are four main components to a plan, as follows.

1. A *trigger event* which may match the top event of an intention.
2. A *prefix* which may match the top of an intention's deed stack.
3. A *guard stack*: the top guard is checked against the agent's state for plan applicability. The rest of the stack is paired off against the rows in the body and may provide additional conditions for that row's execution.
4. A *body* which is the new deed stack that the plan proposes for execution.

Reactive plans do not have trigger events but instead react to the current state (e.g. the beliefs and goals) of the agent. By convention, within the AIL these have a variable

representing an *achieve* goal as a placeholder for a trigger but this is not used by plan selection which focuses simply on checking whether their guard follows from the agent state.

Example Recall our previous example intention:

Event	Deed	Guard	Unifier
`+!clean()`	`+!goto(Room)`	`dirty(Room)`	`Room = room1`
`+!clean()`	`+!vacuum(Room)`	`⊤`	`Room = room1`

A plan that matches this intention is

trigger	`+!clean()`
prefix	`+!goto(Room)`
	`+!vacuum(Room)`
guard	`¬ dirty(Room)`
body	`+!find_dirty_room(Room2)`

So if the current intention was triggered by goal !*clean*() (the trigger event), and it currently intends to go to a room and vacuum it (prefix), *but* that room is not dirty (guard), it proposes, instead, to replace that part of the intention with the goal of locating a dirty room. If this plan was applied, the intention would become:

Event	Deed	Guard	Unifier
`+!clean()`	`+!find_dirty_room(Room2)`	`¬ dirty(Room)`	Room = room1

It is more common for plans to match only intentions which contain unplanned goals (i.e., those associated with the "no plan yet" deed, ε). For instance after a commitment to `goto(Room)` the above intention might appear as:

Event	Deed	Guard	Unifier
`+!goto(Room)`	ε	`⊤`	`Room = room1`
`+!clean()`	`+!goto(Room)`	`dirty(Room)`	`Room = room1`
`+!clean()`	`+!vacuum(Room)`	`⊤`	`Room = room1`

which would match the plan

trigger	`+!goto(Room)`
prefix	ε
guard	`current_floor(ground) ∧ upstairs(Room)`
	`⊤`
body	`+!goto(stairs)`
	`+!goto(Room)`

This would transform the intention to:

Event	Deed	Guard		Unifier
`+!goto(Room)`	`+!goto(stairs)`	`current_floor(ground)`		
		`∧ upstairs(Room)`		`Room = room1`
`+!goto(Room)`	`+!goto(Room)`	`⊤`		`Room = room1`
`+!clean()`	`+!goto(Room)`	`dirty(Room)`		`Room = room1`
`+!clean()`	`+!vacuum(Room)`	`⊤`		`Room = room1`

Applicable plans   Applicable Plans represent an interim data structure that describes how a plan from an agent's plan library changes the current intentions. Essentially an applicable plan states how many rows are to be dropped from the intention and what new rows are to be added. The new rows are generated from an event, a guard stack, a unifier and a stack of deeds. The guard and deed stacks are the same size. The new intention rows are generated by creating a row for each deed and guard on the two stacks and associating the event and unifier with each of those rows (so the event and unifier are duplicated several times).

Therefore, an applicable plan is a tuple, $\langle p_e, p_g, p_{ds}, p_\theta, n \rangle$, of an event p_e, a guard stack p_g, a deed stack p_{ds}, a unifier p_θ and the number of rows to be dropped, n. The applicable plan in the first example above would be

$$\langle +!clean(), [\neg dirty(Room)], [+!find\_dirty\_room(Room2)], \{Room = room1\}, 2 \rangle \tag{1}$$

or in the previous tabular presentation we could represent this as:

Drop 2 Lines
replace with

Event	Deed	Guard	Unifier
`+!clean()`	`+!find_dirty_room(Room2)`	`¬ dirty(Room)`	`Room = room1`

The applicable plan for the second example would be

$$\langle +!goto(Room), [current\_floor(ground) \wedge upstairs(Room); \top],$$
$$[+!goto(stairs); +!goto(Room)], \{Room = room1\}, 1 \rangle \tag{2}$$

or

Drop 1 Line
replace with

Event	Deed	Guard		Unifier
`+!goto(Room)`	`+!goto(stairs)`	`current_floor(ground) ∧`		
		`upstairs(Room)`		`Room = room1`
`+!goto(Room)`	`+!goto(Roo)`	`⊤`		`Room = room1`

Applicable plans are used because many APL reasoning cycles first go through a phase where they determine a list of applicable plans and then move to a phase where they pick one plan to be applied. The function $\mathscr{S}_{\text{plan}}$ (implemented as a Java method

in AIL) is used to select *one* applicable plan from a set. By default, this treats the set as a list and picks the first plan, but it may be overridden by specific languages and applications.

Applicable plan generation method For the purposes of this discussion, we will write plans using the syntax $p_e : p_p : p_{gu} : p_d$ where p_e is the trigger event, p_p is the prefix, p_{gu} is the guard, and p_d is the body.

AIL provides a default function, **appPlans**, for the generation of applicable plans from the current intention and an agent's internal state. This function creates two sets of applicable plans. The first set is, essentially, the applicable plans for continuing to process intention i without any changes:

$$\{\langle \text{hd}_e(i), \text{hd}_g(i), \text{hd}_d(i), \theta \cup \theta^{\text{hd}(i)}, 1\rangle \mid (ag \models \text{hd}_g \theta^{\text{hd}(i)}, \theta) \wedge (\text{hd}_d(i)\theta\theta^{\text{hd}(i)} \neq \varepsilon)\} \tag{3}$$

Here, $\text{hd}_e(i)$ is the top event in i, $\text{hd}_g(i)$ is the top guard, $\text{hd}_d(i)$ is the top deed, and $\theta^{\text{hd}(i)}$ is the top unifier. The notation $ag \models g, \theta$ means that the guard, g, is satisfied by agent ag given unifier θ. The notation $t\theta$ indicates the application of unifier θ to term t. So, for instance, $\text{hd}_d(i)\theta\theta^{\text{hd}(i)}$ is the result of applying the unifiers θ and $\theta^{\text{hd}(i)}$ to the top deed on the intention.

Note that the above generates an empty set if the intention's top deed is the "no plan yet" deed, ε.

The second set is

$$\{\langle p_e, p_{gu}, p_d, \theta^{\text{hd}(i)} \cup \theta, \#p_p\rangle \mid p_e : p_p : p_{gu} : p_d \in P \wedge$$
$$\#p_p > 0 \rightarrow (\text{hd}_e(i) \models p_e \wedge \textit{unifier}(p_p, i) = \theta_e) \wedge ag \models \text{hd}(p_{gu})\theta_e, \theta\}$$

Here, $\text{hd}_e(i) \models p_e$ means that the plan's trigger event follows from the top event on the current intention. This allows for Prolog-style reasoning on plan triggers and is a version of the logical consequence method. Function $\textit{unifier}(p_p, i)$ generates the unifier of the plan prefix with the top n rows of intention i where $n = \#p_p$ (the size of the stack of the plan's prefix).

This general mechanism for deriving applicable plans has proved sufficient for all the APLs implemented in the AIL to date.

3.1.9 Actions

Actions are the means by which an agent affects the external world, and are normally represented as first order terms. In general, when an agent encounters an action, it will cause the execution of "native" code—i.e., code typically developed with traditional programming paradigms rather than code developed for rational behaviour in autonomous systems—to take place in the environment. For instance an action, in a robot, to pick up an object is likely to execute detailed control system code to handle the actual task of picking something up, and this is normally programmed in standard programming languages for control.

3.1.10 Inbox and outbox

Agents maintain *Inboxes* and *Outboxes* for storing messages. The default logical consequence method checks the Outbox (to check that a message has been sent) but not the Inbox—so far the operational semantics implemented have used the Inbox only for temporary storage of messages which are then processed by the agent, but in principle the method could be extended to check for received messages as well. It should be noted that storing *all* messages sent or received is potentially inefficient and languages are not required to use the inboxes and outboxes in the agent class.

3.1.11 Reasoning cycle

Agents are assigned a "reasoning cycle" by the language in which they are written. Each stage of a language's reasoning cycle is typically formalised as a disjunction of semantic rules which define how an agent's state may change during the execution of that stage. Any rule to be used as part of an AIL agent reasoning cycle has to implement a particular Java interface. This can be achieved either by implementing the interface directly or by sub-classing an existing rule. The combined rules of the various stages of the reasoning cycle define the operational semantics of that language. The construction of an interpreter for a language involves the implementation of these rules (which in some cases may already exist in the language's original toolkit) and the implementation of a reasoning cycle, by organising the rules into (the stages of) such a cycle.

In this way, we have implemented, for example, both GOAL (de Boer et al. 2007) and SAAPL (Simple Abstract Agent Programming Language) (Winikoff 2007) interpreters, following their respective operational semantics (Dennis and Fisher 2008) as well as a GWENDOLEN interpreter (Dennis and Farwer 2008).[3] The implementations of these interpreters make use of the AIL operations together with some additional classes specifically added to reproduce faithfully the semantics of those languages.

3.1.12 Formal definition of an AIL agent

An AIL agent can be formalised as a tuple consisting of

- The agent's name (a string), *ag*.
- The agent's belief base (a set of beliefs), *B*.
- The agent's rule base (a set of prolog-style rules for reasoning about guards), *R*.
- The agent's plan library, *P*.
- The agent's current intention (which may be empty), *i*.
- The agent's other intentions, *I*.
- The agent's currently applicable plans, **appPlans**.
- The agent's inbox of unprocessed messages, *In*.
- The agent's outbox of sent messages, *Out*.
- The current stage of the agent's reasoning cycle, *S*.

---

[3]The GWENDOLEN language was developed as a side effect of some of the initial design work on the AIL.

Operational semantic rules operate on this tuple to change the agent's state. It should be noted that individual language implementations may add further elements to this tuple as well as ignore some elements.

A multi-agent system therefore consists of a tuple of several agents and an environment (ξ). The environment has to provide certain services to the agent (the ability to access perceptions, and take actions) but has no formal semantics of its own governing how it should alter from state to state; a model of the environment is assumed to be provided by AIL users in Java.

3.2 Example: Implementing an interpreter for an agent programming language

GOAL (Hindriks et al. 2001; de Boer et al. 2007) is a BDI language introduced by Hindriks et al. to illustrate the use of purely declarative goals in agent programming. An agent is defined by its mental state comprising two sets of formulae: Σ for the agent's beliefs; and Γ for the agent's goals. GOAL assumes an underlying logic on its formula language, $\mathscr{L}$, with an entailment relation $\models_C$; its semantics then defines entailment for mental states as follows:

Definition 1 Let $\langle \Sigma, \Gamma \rangle$ be a mental state:

$$
\begin{aligned}
\langle \Sigma, \Gamma \rangle &\models_M \mathbf{B}\phi && \text{iff } \Sigma \models_C \phi, \\
\langle \Sigma, \Gamma \rangle &\models_M \mathbf{G}\psi && \text{iff } \psi \in \Gamma, \\
\langle \Sigma, \Gamma \rangle &\models_M \neg\phi && \text{iff } \langle \Sigma, \Gamma \rangle \not\models_M \phi, \\
\langle \Sigma, \Gamma \rangle &\models_M \phi_1 \wedge \phi_2 && \text{iff } \langle \Sigma, \Gamma \rangle \models_M \phi_1 \text{ and } \langle \Sigma, \Gamma \rangle \models_M \phi_2.
\end{aligned}
$$

An agent's behaviour is governed by its *capabilities* and *conditional actions*.

Capabilities are associated with a partial function $\mathscr{T} : Bcap \times \wp(\mathscr{L}) \to \wp(\mathscr{L})$, which operates on the belief base Σ in order to alter it. Capabilities may be *enabled* or not for an agent in a particular configuration. If the capability is *not* enabled then $\mathscr{T}$ is undefined. $\mathscr{T}$ is used by the mental state transformation function $\mathscr{M}$ to alter the agent state as follows:

Definition 2 Let $\langle \Sigma, \Gamma \rangle$ be a mental state, and $\mathscr{T}$ be a partial function that associates belief updates with agent capabilities. Then the partial function $\mathscr{M}$ is defined by:

$$
\mathscr{M}(\mathbf{a}, \langle \Sigma, \Gamma \rangle)
$$

$$
= \begin{cases} \langle \mathscr{T}(\mathbf{a}, \Sigma), \Gamma \backslash \{\psi \in \Gamma \mid \mathscr{T}(a, \Sigma) \models_C \psi\} \rangle & \text{if } \mathscr{T}(\mathbf{a}, \Sigma) \text{ is defined,} \\ \text{is undefined for } \mathbf{a} \in Bcap & \text{if } \mathscr{T}(\mathbf{a}, \Sigma) \text{ is undefined} \end{cases} \tag{4}
$$

$$
\mathscr{M}(\mathbf{drop}(\phi), \langle \Sigma, \Gamma \rangle) = \langle \Sigma, \Gamma \backslash \{\psi \in \Gamma \mid \psi \models_C \phi\} \rangle \tag{5}
$$

$$
\mathscr{M}(\mathbf{adopt}(\phi), \langle \Sigma, \Gamma \rangle)
$$

$$
= \begin{cases} \langle \Sigma, \Gamma \cup \{\phi' \mid \Sigma \not\models_C \phi', \models_C \phi \to \phi'\} \rangle & \text{if } \not\models_C \neg\phi \text{ and } \Sigma \not\models_C \phi \\ \text{is undefined} & \text{if } \Sigma \models_C \phi \text{ or } \models_C \neg\phi \end{cases} \tag{6}
$$

Lastly, an agent has a set of conditional actions, Π. Each conditional action consists of a guard, ϕ, and a capability or instruction, **a**, written $\phi \rhd \mathbf{do}(\mathbf{a})$. The conditional actions together with a *commitment strategy* provide a mechanism for selecting which capability to apply next.

Definition 3 Let $\langle \Sigma, \Gamma \rangle$ be a mental state with $b = \phi \rhd \mathbf{do}(\mathbf{a}) \in \Pi$. Then, as a rule, we have: If

1. the mental condition ϕ holds in $\langle \Sigma, \Gamma \rangle$, i.e. $\langle \Sigma, \Gamma \rangle \models_M \phi$, and
2. **a** is enabled in $\langle \Sigma, \Gamma \rangle$, i.e., $\mathscr{M}(\mathbf{a}, \langle \Sigma, \Gamma \rangle)$ is defined,

then $\langle \Sigma, \Gamma \rangle \underset{b}{\longrightarrow} \mathscr{M}(\mathbf{a}, \langle \Sigma, \Gamma \rangle)$ is a possible computation step. The relation $\longrightarrow$ is the smallest relation closed under this rule.

The commitment strategy determines how conditional actions are selected when several potentially apply and this is not specified directly by the GOAL semantics.

3.2.1 GOAL *implemented with AIL*

We discuss the implementation of GOAL with AIL. This work was reported in Dennis and Fisher (2008) so we here focus on some key aspects, in order to use it as an example.

GOAL agents were implemented as sub-classes of the AIL agent class. To model GOAL's mental states we treated the AIL belief base as the GOAL belief base Σ. Although the AIL agent does not contain a specific goal base, we were able to use the method for extracting goals from intentions to represent Γ; a slight modification was needed to the method to return only achievement goals.[4]

The implementation of $\models_M$ was straightforward. Belief formulae of the form $\mathbf{B}(\phi)$ are equivalent to the expressions appearing in AIL's guards and use of AIL's own logical consequence method was sufficient.

3.2.2 *Capabilities and conditional actions*

We chose to model both capabilities and conditional actions as AIL plans, since both control how an agent reacts to its environment.

Inherent in the description of a capability is the idea that the agent performs an action associated with the capability and then makes a number of belief updates. We treated capabilities as *perform goals* because they function as steps/sub-goals that an agent *should* perform, yet they are not declarative. AIL requires the execution of actions to be triggered explicitly so we decided to treat $\mathscr{T}(\mathbf{a}, \Sigma)$ as a function on the belief base paired with an optional action. We write this as $\mathscr{T}(\mathbf{a}, \Sigma) = \mathbf{do}(a) + f(\Sigma)$ and represent it in the AIL as a plan, where the range of f is a deed stack of belief updates. The enabledness of a capability is then governed by the plan guard.

---

[4]As we discuss later, we also used perform goals to model capabilities and did not want these to be returned as part of Γ.

For simplicity, we abstract the full agent state and consider it just as a short tuple of the agent name, current intention, and other intentions: $\langle ag, i, I \rangle$. Conditional actions were modelled as reactive plans with ϕ as the plan guard. In both cases we were able to use the default mechanisms for plan matching and application.

Although we used the default **appPlans** algorithm, we incorporated custom operational semantic rules within this. For instance, for applying a capability we used the rule:

$$\Delta = \{\langle \mathbf{a}, a'; f'(\Sigma) \rangle \mid \mathbf{a} \in Bcap \wedge enabled(\mathbf{a}) \wedge \mathscr{T}(\mathbf{a}, \Sigma) = \mathbf{do}(a') + f'(\Sigma)\}$$

$$\frac{\Delta \neq \emptyset \quad \mathscr{S}_{\text{plan}}(\Delta) = \langle \mathbf{a}, a; f(\Sigma) \rangle}{\langle ag, (\mathbf{a}, \varepsilon); i, I \rangle \rightarrow \langle ag, (\mathbf{a}, a; f(\Sigma)); i', I \backslash \{i'\} \cup \{i\} \rangle} \tag{7}$$

This was a new rule but made use of pre-existing AIL operations (such as $\mathscr{S}_{\text{plan}}$). Δ represents the output of **appPlans**.

By way of illustration, we show some of the Java code for the new "Plan with Capabilities" rule. The key methods of an AIL operational semantics rule are `check-Preconditions` and `apply` which, broadly speaking, represent the conditions for valid rule application and the effects of applying the rule. The precondition for the rule application is that

$$\Delta \neq \emptyset \tag{8}$$

The equivalent code is:

```
public boolean checkPreconditions(AILAgent a)
{
        Intention i = a.getIntention();
        if (i.empty())
                return false;

        if (i.hdE().getGoal().getGoalType() == Goal.performGoal)
        {
                if (a.filterPlans(a.appPlans(i)).isEmpty())
                        return false;
                else
                        return true;
        }

        return true;
}
```

`appPlans` is the Java implementation of **appPlans** and returns the set Δ; the above method returns true if this set is non-empty. `filterPlans` is an overridable method allowing the applicable plan list to be modified in an application specific fashion if desired. By default it does nothing. The test

```
i.hdE().getGoal().getGoalType() == Goal.performGoal
```

restricts the method so it only succeeds if the head event of intention `i` is a "perform goal" (since capabilities are represented by perform goals). The application of the rule is then governed by the code:

```
public void apply(AILAgent a)
{
        Intention i = a.getIntention();
        LinkedList<ApplicablePlan> delta = a.filterPlans(a.appPlans(i));

        GOALAgent ga = (GOALAgent) a;
        ApplicablePlan p = ga.selectPlan(delta, i);

        i.dropP(p.getN());
        i.iConcat(p.getEvent(), p.getPrefix(),
                p.getGuard(), p.getUnifier().clone());
}
```

Here we select a capability, p, from Δ (selectPlan is the implementation of $\mathscr{S}_{plan}$). We drop (using dropP) the appropriate number of lines (p.getN()) from the intention (as specified by the applicable plan), and then add (using iConcat) a new row onto the intention consisting of the event, prefix, guard, and unifier of p. All these methods are provided by the AIL.

3.2.3 The reasoning cycle

AIL assumes a reasoning cycle that passes through a number of explicit stages to which rules are assigned. We therefore analysed the GOAL operational semantics to identify stages. It first selects a conditional action (this involved a new rule, similar to the rule for planning with a capability shown above); then processes the conditional action by applying $\mathscr{M}$; this generally involves selecting a capability and applying it (equivalent to using our new rule above to select a capability) and then processing the effects of that capability as specified by $\mathscr{T}$ (this was easily implemented using existing rules for performing actions and handling belief updates).

3.2.4 Faithfulness of the implementation

Any claim to have implemented the operational semantics of a language is faced with correctness issues involved in transferring a transition system to, in this case, a set of Java classes. Verifying implementations is a complex undertaking. Such a verification effort would be a significant task and falls outside the scope of this work. However, that aside, it is also the case that we have not directly implemented the system presented in de Boer et al. (2007) but a variant of it using, for instance, our representation of capabilities. The question arises: "Are these two transition systems equivalent?" Although we have not done such proofs, we believe they would not represent an insurmountable amount of work.

4 AJPF: verification for multi-agent systems

4.1 Java PathFinder

Java PathFinder (JPF) is an explicit-state Open Source model checker for Java programs (Visser et al. 2003; Java PathFinder 2009). JPF is implemented in Java and provides a special Java Virtual Machine (JVM) running on top of the host JVM

and exploring all executions that a given program can have, resulting from thread interleavings and non-deterministic choices. JPF's backtracking JVM executes Java bytecodes using a special representation of JVM states. Essentially, JPF uses Java's *listener* mechanism to provide a representation of an automaton that is attempting to build a model based on the program execution. As the program proceeds, the listener recognises state changes and checks against user specified properties. At appropriate times this JVM might backtrack and the listener might be reset.

JPF employs state-of-the-art state-space reduction techniques, such as on-the-fly partial-order reduction, i.e., combining instruction sequences that only have effects inside a single thread and executing them as a single transition. This ensures that—out of the box—JPF is a usable (though not very efficient) program model checker; it also should be noted that recent improvements in JPF make it orders of magnitude faster than previous versions. Nevertheless, we have to ensure that the state space *relevant to an agent system* remains as small as possible. For instance, we want to ensure that only relevant backtrack points are stored (i.e., backtracks points in the execution of the operational semantics and not incidental backtrack points in the low-level Java code), thereby limiting the state space and improving the efficiency of model checking.

As well as its usability, we have chosen JPF for several other reasons. Firstly, a large number of APL interpreters are implemented in Java, so a model checker for Java programs was naturally the first choice. Secondly, in order to customise the model checker for our needs and provide it to the public as an Open Source project, we needed the underlying model checker to be Open Source too. Since we were to do significant work in extending the model checker to an agent context (for example, using interfaces), we preferred a model checker working within a standard but higher level programming language. Lastly, we preferred a model checker with an active development team and one that had been developed to provide good extensibility mechanisms. Taking into account these prerequisites, the obvious choice was to base our framework on the Open Source Java model checker JPF.

4.2 AJPF architecture

The AJPF architecture consists of the JPF model checker, the AJPF interfaces, a property specification language (PSL), the controller and AJPF listener classes, and a family of language parsers and translators. A schematic diagram of the architecture was shown in Fig. 1.

A given program, written in one of the supported APLs, is translated into its AIL representation and embedded in an AJPF controller object. The result is a Java program that includes parts of the AIL libraries. This Java program becomes half of the JPF verification target, but instead of invoking standard JPF, we use the JPF extension mechanisms to configure our own *listener*. This implements the property specification language and performs the checking of the property that is supplied with the program. The property is represented as a Büchi Automaton (Sistla et al. 1987). The AJPF controller object combines this Büchi Automaton and the AIL program and then the controller governs when properties are checked and the Büchi Automaton advanced. This is linked to the progress of the agent reasoning cycle.

4.3 The AJPF agent system interfaces—use and semantics

AJPF (see Fig. 1) provides interfaces for model checking a multi-agent program against a property specification written in the Property Specification Language (PSL) introduced later (in Sect. 4.4). AJPF requires that, for any given APL, two Java interfaces are implemented: one for individual agents and another for the overall multi-agent system. This software layer provides an AJPF controller which requests a list of agents from the multi-agent system and encapsulates each of them in a special thread object which alternately calls one reasoning step of the agent, anticipated to be one full run of the reasoning cycle, and then checks this against the specification by calling, for instance, methods that implement belief checking as defined by the specific language. It should be noted that this means the notions of belief, goal, etc. defined within the AIL classes are those invoked by the property checking algorithm. So any language that has an AIL-based interpreter inherits this semantics although individual methods can, of course, be over-ridden. For languages without an AIL-based interpreter, these notions have to be formalised in the implementation of the AJPF interface. Properties are also checked when JPF detects that an "end state" is reached (this could indicate a cycle in the states of a run as well as program termination).

4.3.1 BDI languages without formal semantics

The AIL toolkit is designed for the prototyping of languages which have a clearly defined BDI-based operational semantics. The construction of AJPF however does mean that *in principle* it could be used with agent platforms that have no clear semantics but which are, nevertheless, implemented in Java. The JACK framework (Winikoff 2005), for instance, provides a set of Java classes for programming agents which embody BDI concepts such as beliefs, goals, and plans. In theory, therefore, the appropriate modalities of the property specification language could be defined for JACK programs and, assuming a reasoning steps can be identified, then JACK programs could be model checked using AJPF.

In reality such an attempt is likely to encounter efficiency problems since the JACK code will not make use of atomic sections and state matching hints in order to reduce the model checking state space, so even for a small programs full verification may not be possible in practise.

4.3.2 Environmental models

Since we are dealing with Java programs, we generally assume that the environment in which they operate is also a Java program that can be included in the closed system for model checking. Obviously there are situations where this is not the case, for instance if the programs we are checking are intended to run on robots that operate in the real world. In these cases they have to be model checked against an abstraction of the real environment that would have to be, again, implemented in Java. The same is the case for applications which are intended to be distributed across several different machines. AJPF relies upon Java's thread model for asynchronous execution of agents and can only check a single program executing on a single machine. Therefore the effects of a distributed environment would need to be abstracted into a (possibly multi-threaded) single program environment for model checking in AJPF.

4.4 Specification of properties

In our framework, we are interested in verifying simple properties about *goals*, *beliefs*, *actions*, etc. We do not, in this approach, tackle the verification of properties involving nested modalities, such as beliefs of one agent concerning the beliefs of another agent. Hence, we do not implement a property specification language based on a fully expressive logic of beliefs or knowledge.

A typical property of an agent-based protocol specification could be described by a statement such as: "given an agent a with a goal g and a set of current beliefs $\{b_1, \ldots, b_n\}$, will a eventually believe g?"

Properties are specified at the AJPF level. For agents running on an AIL-based interpreter, the semantics of the properties are already specified as part of the AIL toolkit itself. The PSL allows users to refer to agent concepts at a high level, even though JPF carries out model checking at the Java bytecode level.

We use a property specification language based on propositional linear-time temporal logic (LTL) (Emerson 1990) with added modalities for agents' beliefs, goals, etc. It should be noted that, since it was released as Open Source, JPF no longer supports LTL model checking. Our implementation therefore represents a significant addition to JPF. Also, we do not tackle more complex properties such as those involving nested beliefs and so the PSL defined below is relatively "shallow".

The PSL syntax for property formulae ϕ is as follows, where ag is an "agent constant" and f is a ground first-order atomic formula:

$$\phi ::= \mathbf{B}(ag, f) \mid \mathbf{G}(ag, f) \mid \mathbf{A}(ag, f) \mid \mathbf{I}(ag, f) \mid \mathbf{P}(f) \mid \phi \lor \phi \mid \neg \phi \mid \phi \cup \phi \mid \phi \mathsf{R} \phi$$

Intuitively, $\mathbf{B}(ag, f)$ is true if ag believes f to be true, $\mathbf{G}(ag, f)$ is true if ag has a goal to make f true, and so on (with $\mathbf{A}$ representing actions, $\mathbf{I}$ representing intentions, and $\mathbf{P}$ representing percepts, i.e., properties that are true in the environment).

We next examine the specific semantics of property formulae. Consider a program, P, describing a multi-agent system and let *MAS* be the state of the multi-agent system at one point in the run of P. Let $ag \in MAS$ be an agent at this point in the program execution. Then

$$MAS \models_{MC} \mathbf{B}(ag, f) \quad \text{iff } ag \models f$$

where $\models$ is logical consequence as implemented by the agent programming language. The interpretation of $\mathbf{G}(ag, f)$ is given as:

$$MAS \models_{MC} \mathbf{G}(ag, f) \quad \text{iff } f \in ag_G$$

where ag_G is the set of agent goals (as implemented by the APL). The interpretation of $\mathbf{A}(ag, f)$ is:

$$MAS \models_{MC} \mathbf{A}(ag, f)$$

if, and only if, the last action changing the environment was action f taken by agent ag. Similarly, the interpretation of $\mathbf{I}(ag, f)$ is given as:

$$MAS \models_{MC} \mathbf{I}(ag, f)$$

if, and only if, $f \in ag_G$ and there is an *intended means* for f (in AIL this is interpreted as having selected some plan that can achieve f). Finally, the interpretation of $\mathbf{P}(f)$ is given as:

$$MAS \models_{MC} \mathbf{P}(f)$$

if, and only if, f is a percept that holds true in the environment.

The other operators in the AJPF property specification language have standard LTL semantics (Emerson 1990) and are implemented by the AJPF interface. Thus, the classical logic operators are defined by:

$$MAS \models_{MC} \varphi \vee \psi \quad \text{iff } MAS \models_{MC} \varphi \text{ or } MAS \models_{MC} \psi$$
$$MAS \models_{MC} \neg \phi \quad \text{iff } MAS \not\models_{MC} \phi.$$

The temporal formulae apply to runs of the programs in the JPF model checker. A run consists of a (possibly infinite) sequence of program states MAS_i, $i \geq 0$ where MAS_0 is the initial state of the program (note, however, that for model checking the number of *different* states in any run is assumed to be finite). Let P be a multi-agent program, then

$$MAS \models_{MC} \varphi \,\mathsf{U}\, \psi \quad \text{iff in all runs of } P \text{ there exists a state } MAS_j \text{ such that}$$
$$MAS_i \models_{MC} \varphi \text{ for all } 0 \leq i < j \text{ and } MAS_j \models_{MC} \psi$$
$$MAS \models_{MC} \varphi \,\mathsf{R}\, \psi \quad \text{iff either } MAS_i \models_{MC} \varphi \text{ for all } i \text{ or there exists } MAS_j \text{ such}$$
$$\text{that } MAS_i \models_{MC} \varphi \text{ for all } i \in \{0, \dots, j\} \text{ and}$$
$$MAS_j \models_{MC} \varphi \wedge \psi$$

The common temporal operators $\Diamond$ (eventually) and $\Box$ (always) are derivable from U and R (Emerson 1990).

Typically, APLs do not fully implement logics of belief so, as mentioned above, we use shallow modalities which are like special predicates—however, this does not preclude users, when implementing the AJPF interfaces, from developing a more complex belief logic based on their agent state. The implementation of the modalities defines their semantics (e.g., for *belief*) in that specific language. The AIL implements these interfaces and so defines an AIL specific semantics for the property specification language; supported languages that use the AIL must ensure that their AIL-based interpreters are constructed in a way that makes the AIL semantics of the properties consistent with the language's individual semantics for those modalities (otherwise they cannot use the AIL implementation and will need to override it using the AJPF interface).

4.5 Heterogeneous multi-agent systems

As stated above, the AJPF controller object accepts a set of AJPF agent objects and executes steps of their reasoning cycles, followed by property checks. Since it simply accepts objects that satisfy its interfaces, it is agnostic about the actual semantics of their reasoning cycles and does not require each agent to be using the same semantics at all (although it is necessary for all the agents to interact with the same environment). This makes it as simple to encode a *heterogeneous* multi-agent system as it is to encode a system using only a single agent programming language.

In Dennis and Fisher (2008) we investigated a scenario comprising GWENDOLEN, GOAL and SAAPL agents working together. We were able to successfully implement and verify this system within AJPF. In that work, AIL was found to be appropriate for the three languages and we were able to implement interpreters for SAAPL and GOAL with relative ease.[5] The SAAPL interpreter took about a week to implement and debug while the GOAL interpreter took about two weeks. Once correctly implemented, it was simple to incorporate, run, and verify a heterogeneous multi-agent system.

4.6 Benefits of using AIL and AJPF

The benefits of using the AIL are many, with the main incentives being verification via model checking and the support for heterogeneous multi-agent systems. The flexibility of our approach arises from the fact that the agents can be programmed in a variety of agent-oriented programming languages. This unifying approach to model checking and execution of (heterogeneous) agent systems is an important step towards the practical use of verification techniques, which is essential as dependable systems are increasingly required in many areas of applications of agent technology.

Previous approaches to model checking multi-agent programs focused on a specific APL, e.g. AgentSpeak (Bordini et al. 2006). A language-specific translation of multi-agent systems into models written in the input language of exiting model checkers, generating models that were very difficult to understand, had to be developed; it should also be noted that there is an incredible number of different agent programming languages currently in use in the Agents community. Conceiving and implementing such translations is a tedious, complicated (given the restrictions of typical model checker input languages), and error-prone task that is avoided in the present approach by using an intermediate agent representation that is tailored to multi-agent systems. By lifting the implementation effort away from the model checker to AIL, we make it less tedious and error-prone as we provide tools that are a good match to the operational semantics of agent programming languages.

The architecture of the AIL and AJPF is much more flexible than previous approaches to model checking for agent-based systems. Despite the greater flexibility, we have taken precautions in the construction of the architecture and the internal optimisations of AJPF to ensure that it works relatively efficiently.

We have developed AIL so that new APLs can easily be incorporated into our framework. Even without re-programming a language interpreter using the AIL classes, it is possible to integrate agent programs written in a variety of languages into our verification and execution framework by interfacing their interpreters directly to AJPF. Property specification is uniform amongst all languages that use either the AIL data structures or implement their own notions of belief, goal, action, etc.

4.7 Efficiency issues

In general, model checking suffers from what is known as the *state-space explosion problem*, i.e., the problem that the state space for search increases exponentially as

---

[5]An interpreter for GWENDOLEN had been developed while implementing the AIL.

the system increases in size. It is therefore important to make sure that only essential information is stored in the states that constitute the system to be checked. Using JPF means that the state space that is actually checked is the state space of the Java program representing the AIL agent program. In the remainder, when speaking of the *abstract state space*, we refer to the state space of the multi-agent system bar any additions to the state space that might have been introduced by the translation and use of AIL classes. For efficiency it is, of course, desirable to avoid as many additional states as possible that do not add to the overall behaviour, thus being theoretically harmless.

In spite of this, AJPF is *not* fast. Partly, this is because JPF is itself not a particularly fast model checker (though it is both flexible and appropriate) and partly it is because of the non-trivial additional semantic layer that is added in AJPF. Interestingly recent work on the comparison of model checkers for the GOAL language (Jongmans et al. 2010) concluded that all existing approaches to model-checking agent programming languages that were based on pre-existing model checkers suffered from similar time inefficiencies.

In the following sections, we describe our efforts towards dealing with the issues of efficiency.

4.7.1 Atomic execution

We employ *atomic sections* to reduce the state space of executions whose internal states are not relevant to the execution of the multi-agent system. An atomic section excludes all backtracking within that section. We use this in a number of places where thread interaction is not relevant to the agent transition system.

For instance, using an atomic section for the initialisation phase of the agents and of the MAS leads to a significant speed-up. This portion of the code is executed many times as JPF backtracks (thereby initialising the agents in a different order) and substantial savings result from this. Further use of atomic sections in the reasoning cycle also help improve efficiency.

4.7.2 State matching

In model checking it is important to take care that, in the system to be verified, the states that are conceptually identical do not contain any components that would make the model checker distinguish between them. It is therefore essential to restrict the data structures to the abstract agent system, hiding from the model checker any components that might have been introduced for operational reasons or to provide statistics, such as counters, that do not inherently belong to the agent state. JPF's state matching is an important mechanism to avoid unnecessary work. The execution state of a program mainly consists of heap and thread-stack snapshots. During the execution of a program, JPF checks every state it reaches against its history. If an equivalent state had been reached before, there is no need to continue along the current execution path. In this case, JPF backtracks to the last unexplored non-deterministic choice.

JPF provides a means to tag variables or parts of data structures, so that the model checker ignores them. This is essential to get state matching to work in the presence of

counters, etc. JPF supports this abstraction through the `@FilterField` annotation. Applying this to (part of) a data structure explicitly declares the structure to be exempt from state matching.

JPF uses its own internal mechanisms for state matching based on Jenkins hashes (Jenkins 1997). It is outside the scope of this work to discuss the implementation and the trade-offs involved in the efficiency of compiling the hash. We have observed that, since execution of the Java engine is comparatively slow and is independent of the implementation of AIL, there are important efficiency gains to be made if as many states can be matched as possible.

4.7.3 Property checking

Checking temporal properties, as defined in our property specification language, can cause branches of the search space where the property automaton branches. We therefore limit the places in which these properties need to be checked. In general we only check properties at the end of the execution of a whole reasoning cycle rather than, for instance, after the application of every operational semantic rule.

5 Evaluation

Having described both AIL and AJPF, we now provide a number of scenarios showing the whole system in action. While these examples are relatively simple, they exhibit all the functionality of the MCAPL system.

In these examples we target one particular BDI language. Although this language is simple, it is designed to exhibit many features common to BDI languages in general. Agents are represented as sets of initial beliefs and goals together with a library of plans. A multi-agent system is a set of agents, together with an environment, through which communication occurs and in which actions are performed.

Below, we first describe the various scenarios and then, in Sect. 5.2, discuss their verification.

5.1 Verified scenarios

5.1.1 Contract Net example

The *Contract Net* scenario (Smith and Davis 1980) is a well-known, and widely used, model of cooperation in distributed problem-solving. Essentially, a particular agent (the *manager*) broadcasts tasks (goals) to be accomplished, and then agents capable of doing so bid for the contract. In real scenarios, the bidding, allocation, and sub-contracting can be quite complex. However, we consider a very simple version: the manager does not broadcast to all the agents in the system at once but instead contacts them in turn; there is no bidding process nor sub-contracting; agents volunteer for a task if, and only if, they can perform it; and the manager simply accepts the first proposal it receives.

We investigated the model checking of this scenario with up to six agents attempting to achieve either one or two goals requested by the manager. The code for the version with two bidding agents and two goals can be found in Appendix B.2.

5.1.2 Auction scenarios

We also considered a simple auction example. The basic idea of an *auction* (Vickrey 1961; Klemperer 2004) is at the heart of many multi-agent scenarios (Boutilier et al. 1997). Not only are auctions central to e-commerce applications (Collins et al. 2009; Fortnow et al. 2008), they are implicit within many *market-based* approaches to agent computation (Walsh and Wellman 1998). These include areas where resource allocation or task allocation is required, for example in telecommunications (Haque et al. 2005; Gibney et al. 1999), electricity supply management (Corera et al. 1996), agent mobility (Bredin et al. 2003), logistics (Dash et al. 2007), and scheduling (Reeves et al. 2005). However, although much work has been carried out on deep analysis of auction mechanisms, such as through formal mechanism design (Wooldridge et al. 2007), the analysis of *implementations* of auction mechanisms has lagged behind. While there has been some work on the formal verification of auction implementations, such as Doghri (2008), this has lacked an agent perspective. Thus, the more sophisticated agent aspects such as goals, intentions, beliefs, and deliberation are not typically verified within an auction context.

The basic version of this study is initially very simple. We describe the basic scenario below and then, in subsequent sections, we describe more sophisticated variants, each becoming increasingly realistic.

A very basic auction The idea here is simple. A number of agents (in the diagram below, three) make bids of some value to an auctioneer agent. The auctioneer agent then awards the resource to the highest bidder and announces this. This cycle can then repeat, if necessary (note that, in our verified scenarios, the bidding process does *not* cycle).
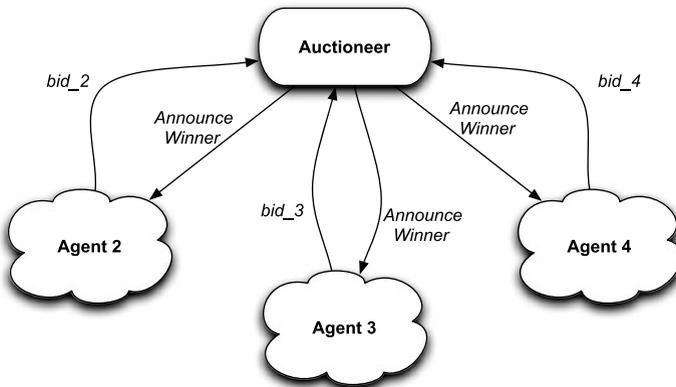
Versions of this scenario with increasing numbers of agents were implemented in GWENDOLEN. The code for the four-agent version can be found in Appendix B.3.

Auction coalition scenario The above basic scenario was next extended to include the possibility of *coalitions* (Sandholm and Lesser 1997; Konishi and Ray 2003). In our model, a coalition occurs when several agents collaborate by pooling their bid

resources in order to win the auction. For example, if three agents x, y, and z bid 100, 150, and 200 respectively, then z ought to win every time. However, if x and y form a coalition, their collective bid of 250 will be enough to win the auction.

A simple coalition scenario was implemented in GWENDOLEN with an auctioneer and a variable number of bidders. A version of the code for this scenario with 4 agents is shown in Appendix B.4. In that version, all but one of the bidders bid straight away, but one of the agents attempts to form a coalition by communicating with one of the other bidders. The contacted bidder agrees to form the coalition, and informs the coalition former of its bidding amount. The coalition instigator then combines its own bidding amount with that of its coalition partner and submits this bid to the auctioneer. Then, having received all of the bids, the auctioneer announces the winner.

Below, Agent 3 instigates the coalition:
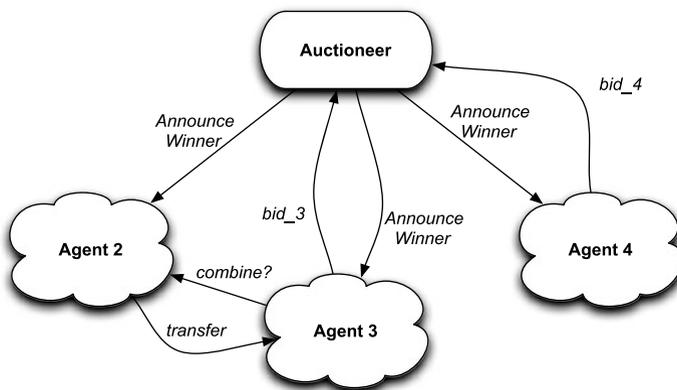
The main difference in the implementation of this scenario, as compared with our earlier one, is that one agent, Agent 3, has a goal to form a coalition. Agent 3 then contacts Agent 2 and proposes a coalition. If Agent 2 agrees then Agent 3 can now bid a winning 250 (i.e., $100 + 150$). Clearly, we would like to verify that this approach does, indeed, lead to Agent 3 winning the auction. This is one of the properties we verify in Sect. 5.2.

Dynamic auction coalition scenario   A further variant on the auction coalition scenario was implemented. In this case, a round of bidding takes place in which all agents bid. Then, after an agent discovers that it has lost the auction, it sends a message to another agent (excluding the previous winner) to form a coalition. Then, the agents bid again. Sample code can be found in Appendix B.5.

Coalition trust scenario   This auction scenario is similar to that described in Sect. 5.1.2, except the coalition forming agent now has a belief about which other agent(s) it can trust, i.e., the other agents with which it would *prefer* to form a coalition. This trust aspect is static, that is, the coalition-forming agent starts the auction with belief(s) about which agents it can trust, and these do not change during the auction. Sample code for this scenario can be seen in Appendix B.6.

Dynamic trust scenario This final auction scenario builds upon the previous one. Here, if the coalition-forming agent loses the auction, it tries to form a coalition with an agent it trusts. Then, if its coalition is successful in winning the auction, it stops. However, if its coalition is *unsuccessful* then it no longer believes that it can trust the other agent in the coalition, and will try to form another coalition with another agent it trusts (excluding the winner). Again, sample code for this scenario can be seen in Appendix B.7.

5.1.3 Trash collection robots

Our last example is based on the garbage collection agents reported in Bordini et al. (2003, 2006). This example was previously written in AgentSpeak and then verified in both the Spin (Holzmann 2003) and JPF model checkers. This work was the immediate precursor of the work reported here that represents an attempt to re-engineer the model checking system to make it more generic. In our previous examples, we were interested in the effects on the system of adding more agents to it. Here, our intention was to gain some idea of the cost of the more generic architecture. The scenario investigated involved two robots (theoretically on Mars) detecting and burning two pieces of garbage placed at random on a 5 by 5 grid. The first robot searches for garbage, picks it up and takes it to the second robot. The second robot then picks up the garbage and incinerates it. We translated the original AgentSpeak code directly into GWENDOLEN. This code is shown in Appendix B.8.

5.2 Results

We investigated a number of aspects of the verification of our case studies, as reported in the sections below.

5.2.1 Effect of scaling the program

We started by investigating the effect of the complexity of the program on the size of the state space. As a crude measure of an increase in the program complexity we investigated the effect of adding an additional agent into the contract net and auction scenarios. We did not investiage the trash collection robot scenario in this way since this scenario involved no communication between the agents and was developed primarily in order to provide a comparison with previous work. In the contract net scenario, we verified that eventually the manager believed all its goals were achieved. In the auction scenario, we verified that eventually the agent making the highest bid believed it had won.

The effects of adding additional agents on the state space are shown in Fig. 2. As can be seen, the size of the state space increases exponentially as more agents are added into a scenario. This represents a typical result for a model checking exercise.

Although less informative since execution time can easily be affected by factors other than the program under consideration, we also investigated the effect of extra agents against the time taken for a program to be verified. The results of this are shown in Fig. 3. As can be seen, although we have comparatively few states in our
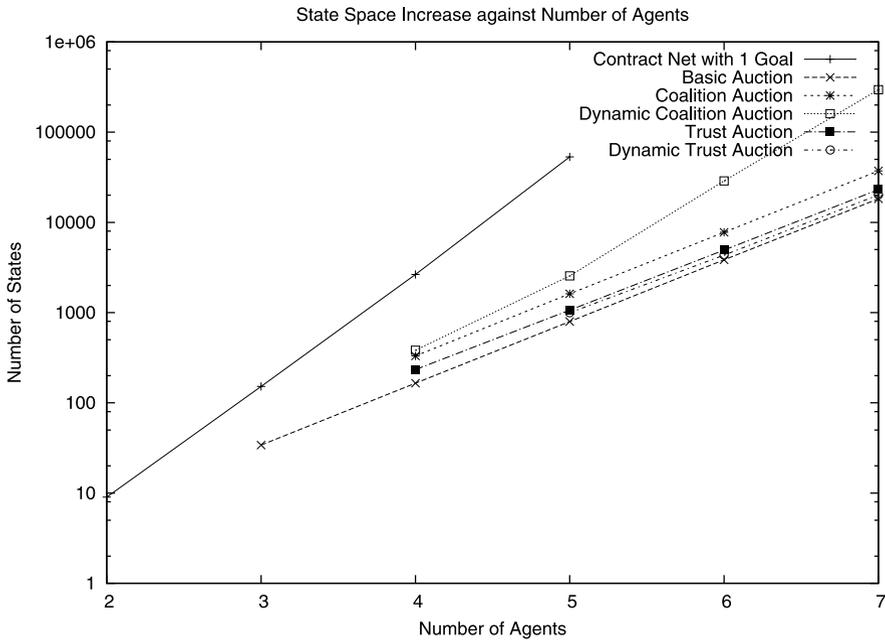
Fig. 2 State space increase as number of agents increases

space compared to many model checking systems, we nevertheless take considerable time to verify a program. This is because each transition in the state graph of the model checker takes significant time to execute as the JPF JVM processes many bytecode instructions.

5.2.2 Effect of the property size

As well as checking properties about the beliefs of the agents, we also checked properties related to their goals, actions, intentions, etc. In theory, since the BDI modalities of the property specification language are treated as propositions by the property automata, the effect of the property on the model checking state space should be the same as for standard LTL model checking (Holzmann 2003).

In the Contract Net scenario, on a system consisting of three agents bidding to perform one goal, we checked the validity of the properties shown in Fig. 4 with their resulting state space size. As predicted, the size of the state space appears primarily to depend upon the LTL elements with disjunctions involving eventualities creating the largest state spaces.

5.2.3 Comparison with previous work

Lastly, we compared the performance of our model checking system against the previous work for model checking AgentSpeak systems. Bordini et al. (2003, 2006) re-

Time Increase against Number of Agents
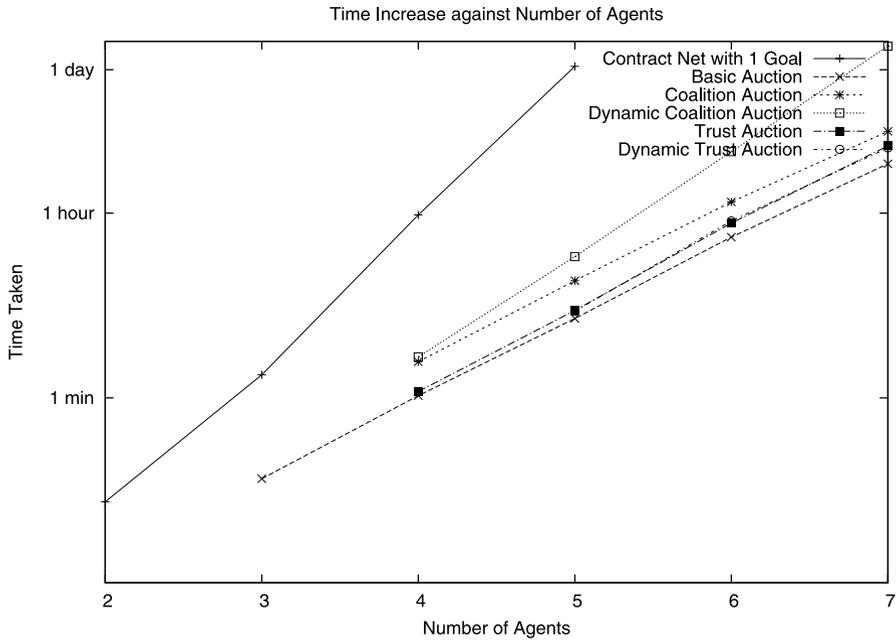
Fig. 3 Time taken as number of agents increases

Property	No. of states
$\Diamond\mathbf{B}(ag1, g)$	152
$\Diamond\mathbf{P}(g)$	106
$\Box(\mathbf{G}(ag2, g) \rightarrow \Diamond\mathbf{A}(ag2, g))$	281
$\Diamond(\mathbf{G}(ag2, g) \vee \mathbf{G}(ag3, g))$	79
$\Diamond(\mathbf{I}(ag2, g) \vee \mathbf{I}(ag3, g))$	79
$\Box(\mathbf{G}(ag3, respond(g, ag1)) \rightarrow \Diamond(\mathbf{B}(ag2, award(g)) \vee \mathbf{A}(ag3, a)))$	294

Fig. 4 Varying the property checked

ported that verifying the property $\Diamond(\mathbf{I}(r1, continue(check)) \wedge \mathbf{B}(r1, checking(slots))$ took 333,413 states and 65 seconds to verify in Spin and 236,946 states and 18 hours to verify in JPF. In our (JPF-based) system the verification of the same property used 23,655 states and took 9 hours. Clearly the Spin based system remains vastly superior in terms of efficiency. It is hard to accurately compare the results for the JPF based systems since JPF itself has been the subject of continuous development. However, it is clear that the generic approach is unlikely to be *worse* in terms of efficiency than the language-specific approach taken previously, although it is known that JPF is currently much faster than it used to be.

6 Conclusions

6.1 Summary

In this paper we have described the development of the MCAPL framework, incorporating the AIL intermediate semantic layer and the AJPF enhanced model checker. We have seen that the AIL semantic structures are sufficiently expressive to allow developers to capture a range of BDI programming languages and that the semantic rules fit well with the modified AJPF model checker. The efficacy and generality of the AIL toolkit has also been established by the implementation of a variety of different agent programming languages and the verification of multi-agent systems implemented in those languages. Interpreters have been implemented for GWENDOLEN (Dennis and Farwer 2008), GOAL (de Boer et al. 2007), SAAPL (Winikoff 2007), and ORWELL (Dastani et al. 2009). Interpreters for AgentSpeak (Bordini et al. 2007) and 3APL (Dastani et al. 2005) are also being developed. Importantly, the MCAPL framework is also appropriate for verifying *heterogeneous* multi-agent systems, as well as homogeneous ones.

The overall approach has been designed and implemented. It has also been tested on some small multi-agent programs: variations of the *contract net* protocol (Smith 1980) and *auction systems*, but with five or fewer agents (Webster et al. 2009). Specifically, we have focused on a series of scenarios of increasing complexity in order to demonstrate that, although the difficulty of the model checking task increases with each scenario, it is nevertheless realistic to model-check the properties of interesting multi-agent implementations within a reasonable time.

Thus, the MCAPL framework provides a generic harness for automatically verifying agent software. Clearly, for bigger scenarios, improved efficiency will be required (see the discussion in the next section), but the examples implemented and verified in this paper demonstrate that simple properties of multi-agent systems can already be tackled.

6.2 Efficiency problems

A typical problem in model checking, particularly of concurrent systems (where various entities have independent, yet interacting, behaviour), is that of *state space explosion*. The model checker needs to build an in-memory model of the states of the system, and the number of such states grows exponentially for example in the number of different entities being modelled.

Even with refined representation techniques, such as the BDDs used in *symbolic model checking* (Burch et al. 1992), the formulae/structures required to represent the state spaces of realistic systems are huge. JPF is an explicit-state, on-the-fly model checker, and a further problem is that the underlying JPF virtual machine is rather slow. Thus, our current verification system is also slow (although recent work shows that its performance is comparable to a similar system implemented in Maude; Jongmans et al. 2010). Although speed is the main problem, space required can also be problematic (Bordini et al. 2008) (though note that the slow examples above actually explore fewer than 500,000 states in total). We should also note that the success of program model checking relies a great deal on *state-space reduction techniques*,

which we have also adapted for agent verification in Bordini et al. (2009), but have not yet implemented to work at the AIL level.

However, our approach is no less efficient than the language-specific work reported in Bordini et al. (2006). Thus, it is our belief that the generic design principles embodied in the MCAPL framework could be transported to other model checking systems and it is not the *generality* of the framework which is the main issue in terms of efficiency.

6.3 Future work

Our proposed future work falls into three main areas.

Firstly, we would like to extend the agent programming languages available within the system to include, at the least, the *Jason* (Bordini et al. 2005b, 2007) implementation of AgentSpeak and 3APL (Hindriks et al. 1998; Dastani et al. 2005). At the same time we would like to improve the support for the languages we already have implemented in terms of supplying more complete parsers and translators for them so that programs written for other implementations of those languages can be easily imported into our system and run.

Secondly, we intend to improve the model checking aspects of the framework. In particular, we would like to investigate the use of "mixed execution" in JPF (D'Amorim 2007). This would allow us to delegate the operation of parts of the Java code to the native, efficient Java Virtual Machine rather than using the JPF virtual machine. This involves identifying appropriate methods and data structures which are irrelevant to the correct storage of the system state for backtracking. We have made some initial, inconclusive, investigations into delegating the unification algorithm in this way, but work elsewhere suggests that we should be able to achieve significant time improvements.

Lastly, we are interested in replacing the JPF back end to the system with a different model checker such as Spin (Holzmann 2003) or NuSMV (Cimatti et al. 2002). Previous work (Bordini et al. 2003, 2006) suggests that a considerable speed up may be possible in another model checker but that more work would be required in creating a framework in which the model checker could simply tackle a range of different agent programming languages. It is our belief that much of the design work reported here could be adapted to an alternative system.

Appendix A: AIL operational semantics rules

A.1 Introduction and notation

The purpose of this chapter is to provide a reference for the presupplied rules that may be used in the operational semantics of a language implemented in the AIL. We have tried to present these in as clear a fashion as possible excluding implementation details where possible.

An agent can be viewed as a large tuple consisting of the fields of the AILAgent class. Writing out every element of this tuple makes the presentation largely unreadable, therefore we restrict ourselves to including only those elements of the tuple that are relevant to the rule itself. These can be identified by the naming conventions shown below and, where there is a possibility of ambiguity we will indicate these with equalities—i.e. $i = (a, \varepsilon)$ indicates that the current intention is (a, ε). Where a value is changed by the transition it will often be primed to indicate the new value—e.g. $i' = null$ shows that the new value of the current intention is *null*.

ag	The name of the agent.
B	The agent's Belief base.
i	The current intention.
I	The agent's other intentions.
Pl	The agent's applicable plans.
A	The agent's queue of pending actions.
In	The agent's inbox.
Out	The agent's outbox.
ξ	The agent's environment.

Many of the operational rules make a check on a deed to see what type it is (e.g. the addition of a belief, the deletion of a goal). We represent these checks implicitly using notation as follows:

a	An AIL data structure of action type.
b	An AIL data structure of belief type.
$+b$	A belief addition.
$-b$	A belief removal.
b{source}	A belief, from source *source*.
$!_\tau g$	A goal of type τ.
$+!_\tau g$	A goal addition.
$-!_\tau g$	A goal drop.
$\times !_\tau g$	A goal there is a problem with.
`lock`	An AIL lock structure.
`unlock`	An AIL unlock structure.
$\uparrow^{ag} m$	A message *m* sent to *ag*.
$\downarrow^{ag} m$	A message *m* received from *ag*.
$\top$	An AIL structure who's logical content is trivially true.
ε	A special marker indicating that some event has no plan yet.

Many of the rules make reference to methods that exist in the AILAgent class. Obviously subclassing these methods potentially changes the semantics of the rule. This is intentionally the case. Table 1 shows the notation used for methods in the AILAgent class.

Rules that specifically deal with accessing information from the environment, ξ, refence methods specified in the AILEnv interface that have to be implemented by any environment. Again this means the semantics of the rules will depend upon the environment used. Table 2 shows the notation used for methods in the AILEnv interface.

Table 1 Notation for methods in the AILAgent class

Notation	Method name	Description
allsuspended	allintentionssuspended	All the intentions in the agent are suspended.
consistent(b)	consistent(b)	b is consistent with the belief base (defaults to true).
appPlans(i)	appPlans(i))	Generate application plans.
filter(Pl)	filterPlans(Pl)	Filter plans from the set (defaults to none).
$ag \models gu, \theta$	$\theta = $ logicalconsequence(gu)	The agent believes the guard gu given unifier θ.
relevant(s, s')	relevant(s, s')	s and s' are sources of information relevant to each other (defaults to true).
$\mathscr{S}_{\text{int}}(I)$	selectIntention(I)	Select an intention from the set I.
$\mathscr{S}_{\text{plan}}(Pl)$	selectPlan(Pl)	Select a plan from the set Pl.

Table 2 Notation for methods in the AILEnv interface

Notation	Method name	Description
$\xi.\textbf{do}(a)$	executeAction	Executing an action in the environment.
$\xi.\textbf{getMessages}()$	getMessages	Returns new messages.
$\xi.\textbf{Percepts}(ag)$	getPercepts	Returns new perceptions.

Many of the rules also manipulate, or check information about single intentions. Again these reference methods in the Intention class. Table 3 shows the notation used for methods in the Intention class.

Lastly we use a few functions as shorthand for more complex processes with the AIL toolkit:

$oldPercepts(P)$	Any beliefs in an agent's belief base which are marked as percepts (i.e. their source is percept) which are not in the set P.
$wake(ag)$	Unset any flags telling the agent to sleep it's thread next opportunity.
$unify(l_1, l_2)$	Unify the two Unifiable structures l_1 and l_2. This generally creates an empty unifier and then calls its unifies method.
$unsuspend(I)$	Unsuspend all the intentions in I.
$\tau_a(a)$	Returns the "type" of action, a. Useful when a semantics wants to separate actions into categories and treat them differently.
$d\theta$	Represents the application of a unifier, θ, to some data structure, d. d may be an action, a belief, a goal, a message, an event, a guard or a deed.

A.2 The rules

1. ApplyApplicablePlans

$$\frac{Pl \neq \emptyset \quad \langle 0, e, g; G, Ds, \theta \rangle = \mathscr{S}_{\text{plan}}(Pl) \quad g \neq \top}{\langle i, Pl \rangle \rightarrow \langle i' = \texttt{new}(+state(g), \top; G, Ds, \theta, \texttt{self}), Pl' = \emptyset \rangle} \quad (9)$$

Table 3 Notation for methods in the Intention class

Notation	Method name	Description
U_θ	compose	Compose the unifier with the top unifier on the intention.
$\mathrm{drop}(g)$	dropGoal	Drop all rows from the intention until one is reached with $+!_\tau g$ as it's event.
$\mathrm{drop}_\mathrm{p}(N, i)$	dropP	Drop N rows from the top of the intention.
$\mathrm{empty}(i)$	empty	The deed stack of the intention is empty.
$\mathrm{events}(i)$	events	The set/stack of events associated with the intention.
$\mathrm{hd}_d(i)$	hdD	The top deed on the intention.
$\mathrm{hd}_e(i)$	hdE	The top event on the intention.
$\mathrm{hd}_g(i)$	hdG	The top guard on the intention.
$\theta^{\mathrm{hd}(i)}$	hdU	The top unifier on the intention.
@	iConcat	Add a new event, deed stack, guard stack and unifier to the top of the intention.
$;_p$	iCons	Add a new event, deed, guard and unifier as the top row of the intention.
$\mathrm{new}(e)$	Intention	Create a new intention from the event e.
$\mathrm{new}(e, source)$	Intention	Create a new intention from the event e and the source *source*.
$\mathrm{new}(e, Gs, Ds, \theta, source)$	Intention	Create a new intention from an event, guard stack, deed stack, unifier and source.
$\mathrm{lock}(i)$	lock	Mark the intention as locked.
$\mathrm{locked}(i)$	locked	The intention is locked.
$\mathrm{noplan}(i)$	noplan	The intention has not been planned (i.e. the deed stack is empty or contains only the ε "no plan yet" deed).
$\mathrm{suspend}(i)$	suspend	Mark the intention as suspended.
$\mathrm{tl}_i(i)$	tlI	Drop the top deed (with associated event, guard and unifier) from the intention.
$\mathrm{unlock}(i)$	unlock	Mark the intention as unlocked.

$$\frac{Pl \neq \emptyset \quad \langle N, e, g; G, Ds, \theta \rangle = \mathscr{S}_{\mathrm{plan}}(Pl) \quad N > 0 \vee g = \top}{\langle i, Pl \rangle \rightarrow \langle i' = (e, g; G, Ds, \theta) @ \mathrm{drop}_\mathrm{p}(N, i), Pl' = \emptyset \rangle} \tag{10}$$

Notes: This rule selects a plan from the agent's applicable plans. The plan is represented as a tuple of the number of rows to be dropped, the trigger event, the plan's guard stack, deed stack and unifier.

If it is a reactive plan then N is equal to 0. In this case a new trigger is created $+state(g)$ where g is the top guard on the plan's guardstack. This is supposed to represent the state of the world that triggered the plan. A new intention is created from the applicable plan.

Otherwise the applicable plan is "glued" to the top of the current intention.

2. DirectPerception

$$\frac{P = \xi.\textbf{Percepts}(ag)}{\langle ag, B, In \rangle \to}$$
$$wake(\langle ag, B' = B \cup P \setminus oldPercepts(P), In' = In \cup \xi.\textbf{getMessages}()\rangle) \quad (11)$$

Notes: A simple perception rule. It adds all percepts to the belief base and removes all beliefs no longer perceived. It also add all messages to the inbox. A key part of the working of the rule depends on AIL's annotation of all beliefs in the belief base with a source and its use of a special annotation for beliefs whose source is perception.

3. DirectPerceptionwEvent

$$\frac{\begin{array}{cc} P = \xi.\textbf{Percepts}(ag) & P^- = oldPercepts(P) \\ I_1 = \{\texttt{new}(+b)|b \in P\} & I_2 = \{\texttt{new}(-b)|b \in P^-\} \end{array}}{\langle ag, B, I, In \rangle \to}$$
$$wake(\langle ag, I' = I \cup I_1 \cup I_2, B' = B \cup P \setminus P^-, In' = In \cup \xi.\textbf{getMessages}()\rangle) \quad (12)$$

Notes: Similar to DirectPerception this rule also creates new intentions triggered by the addition (or removal) of all the beliefs allowing the agent to react to the changes.

4. DoNothing

$$\frac{}{A \to A} \quad (13)$$

Notes: The `DoNothing` rule, as its name suggests, makes no changes to the state of the agent. This is intended as a default rule that can be used in a reasoning cycle stage to do nothing if none of the other rules in the stage apply but nevertheless still allow the stage to progress.

5. DropIntention

$$\frac{I \neq \emptyset \quad \mathscr{S}_{\text{int}}(I) = (i', I')}{\langle ag, i, I \rangle \to \langle ag, i', I' \rangle} \quad (14)$$

Notes: `DropIntention` is really intended for sub-classing. It simply drops the current intention, i, and selects a new one from the intention set. A subclass would be expected to place extra conditions on i to make sure they are only dropped in very specific circumstances. See `DropIntentionIfEmpty` for an example.

6. DropIntentionIfEmpty

$$\frac{i \neq null \quad \texttt{empty}(i) \quad I \neq \emptyset \quad \mathscr{S}_{\text{int}}I = (i', I')}{\langle ag, i, I \rangle \to \langle ag, i', I' \rangle} \quad (15)$$

Notes: `DropIntentionIfEmpty` drops the intention i if it is empty and selects a new current intention from the intention set. The additional $i \neq null$ is necessary since a few rules can leave the agent state with no current intention.

7. GenerateApplicablePlans

$$\overline{\langle ag, i, Pl \rangle \rightarrow \langle ag, i, Pl' = \textbf{filter}(\textbf{appPlans}(i)) \rangle} \tag{16}$$

Notes: This rule considers all the plans in the agent's plan library and examines all possible instantions of these plans, if there is more than one. It converts these instantiated plans to Applicable Plans, filters them according to any language specific heuristics (as defined by over-riding of the filterPlans method), and places them in the agents Applicable Plan list. The rule is primarily indented for subclassing by more sophisticated rules.

8. GenerateApplicablePlansEmpty

$$\frac{\textbf{filter}(\textbf{appPlans}(i)) = \emptyset \quad \text{noplan}(i)}{\langle ag, i, Pl \rangle \rightarrow \langle ag, i, Pl' = [(1, \text{hd}_e(i), [], [], \emptyset)] \rangle} \tag{17}$$

$$\frac{\textbf{filter}(\textbf{appPlans}(i)) = \emptyset \quad \neg\text{noplan}(i)}{\langle ag, i, Pl \rangle \rightarrow \langle ag, i, Pl' = [(0, \text{hd}_e(i), [], [], \emptyset)] \rangle} \tag{18}$$

Notes: This is a special case of the GenerateApplicablePlans rule for when the set of applicable plans is empty. It does two different things depending on whether or not the current intention has a plan at the top. If it does then the rule provides an applicable plan that will leave the intention unchanged allowing the plan to continue processing. If not it provides a plan that will drop the top row of the intention (for instance if the intention indicates a belief change event then the absence of a plan means that the system has no need to respond to that event. This rule will thus cause that belief change notification to be dropped).

9. GenerateApplicablePlansEmptyProblemGoal

$$\frac{\textbf{filter}(\textbf{appPlans}(i)) = \emptyset \quad \text{noplan}(i) \quad \text{hd}_e(i) = +!_\tau g}{\langle ag, i, Pl \rangle \rightarrow \langle ag, i, Pl' = [0, \text{x}!_\tau g, [\varepsilon], [\top], \theta^{\text{hd}(()i)})] \rangle} \tag{19}$$

$$\frac{\textbf{filter}(\textbf{appPlans}(i)) = \emptyset \quad \text{noplan}(i) \quad \neg\text{hd}_e(i) = g}{\langle ag, i, Pl \rangle \rightarrow \langle ag, i, Pl' = [(1, \text{hd}_e(i), [], [], \emptyset)] \rangle} \tag{20}$$

$$\frac{\textbf{filter}(\textbf{appPlans}(i)) = \emptyset \quad \neg\text{noplan}(i) \quad \neg\text{hd}_e(i) = g}{\langle ag, i, Pl \rangle \rightarrow \langle ag, i, Pl' = [(0, \text{hd}_e(i), [], [], \emptyset)] \rangle} \tag{21}$$

Notes: This is a further specialisation of GenerateApplicablePlansEmpty. In this case if the current intention has no plan and the trigger event at the top of the intention is a goal then, instead of simply dropping that row (as GenerateApplicablePlansEmpty does) it adds a new row triggered by a "problem goal" event. The agent can then react to that problem goal if it has an appropriate plan.

10. GenerateApplicablePlansIfNonEmpty

$$\frac{\textbf{filter}(\textbf{appPlans}(i)) \neq \emptyset}{\langle ag, i, Pl \rangle \rightarrow \langle ag, i, Pl' = \textbf{filter}(\textbf{appPlans}(i)) \rangle} \tag{22}$$

Notes: As `GenerateApplicablePlans` except with a check for non-emptiness.

11. HandleAction

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \xi.\mathbf{do}(a\theta_b) = \theta_a}{\langle ag, i \rangle \rightarrow \langle ag, i' = \mathrm{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_a)\rangle} \tag{23}$$

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \neg\xi.\mathbf{do}(a\theta_b)}{\langle ag, i \rangle \rightarrow \langle ag, i' = \mathrm{tl}_i(i)\mathtt{U}_\theta\theta^{\mathrm{hd}(i)}\rangle} \tag{24}$$

Notes: This is a basic action handling rule. It attempts to execute the action in the environment (**do**). If this succeeds it gets a unifier which is handed back to the intention. If the action fails it is simply ignored. Most languages will want to explicitly handle action failure in some way, possibly by sub-classing this rule.

12. HandleActionwProblem

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \xi.\mathbf{do}(a\theta_b) = \theta_a}{\langle ag, i \rangle \rightarrow \langle ag, i' = \mathrm{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_a)\rangle} \tag{25}$$

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \neg\xi.\mathbf{do}(a\theta_b) \quad \mathrm{hd}_e(i) = +!_{\tau_g}g}{\langle ag, i \rangle \rightarrow \langle ag, i' = (\mathrm{x}!_{\tau_g}g, \theta^{\mathrm{hd}(i)} \cup \theta_a);_p i\rangle} \tag{26}$$

Notes: This extends `HandleAction` with some failure handling. If the action appears on the deed stack because of a goal commitment then the intention gets a new intention noting there is a problem with the goal. The agent can then react to this—e.g. by attempting the action again or dropping the goal or in some other fashion.

13. HandleAddAchieveTestGoal

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = +!_{\tau_g}g \quad \tau_g = a \vee \tau_g = t \quad ag \models g\theta_b, \theta_g}{\langle ag, i \rangle \rightarrow \langle ag, i' = \mathrm{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_a \cup \theta_g)\rangle} \tag{27}$$

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = +!_{\tau_g}g \quad \tau_g = a \vee \tau_g = t \quad ag \not\models g\theta_b}{\langle ag, i \rangle \rightarrow \langle ag, i' = (g\tau_g, \theta^{\mathrm{hd}(\cup)}\theta_b);_p i\rangle} \tag{28}$$

Notes: This rule handles the commitment to an achieve or test goal. These are detected in the condition $\tau_g = a \vee \tau_g = t$. An achieve or test goal is one that triggers a plan if it not already believed but does no more than set a unifier if it is. If it is to trigger a plan then we register the commitment to planning the goal as an event on the top of the intention stack.

14. HandleAddBelief

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = +b \quad \mathbf{consistent}(B, b)}{\langle ag, B, i, I \rangle \rightarrow}$$
$$\langle ag, B' = B \cup b\{\mathtt{src}(i)\}, i' = \mathrm{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), I' = \mathit{unsuspend}(I)\rangle \tag{29}$$

Notes: A basic rule for adding a new belief to the belief base. It assigns a source to the belief, which is the source of the intention. As a side effect it "unsuspends" all intentions.

15. HandleAddBeliefwEvent

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = +b \quad \textbf{consistent}(B, b)}{\begin{array}{c} \langle ag, B, i, I \rangle \rightarrow \\ \langle ag, B' = B \cup b\{\mathrm{src}(i)\}, \\ i' = \mathrm{tl}_i(i)\mathrm{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), I' = \mathit{unsuspend}(I) \cup \mathrm{new}(+b, \mathrm{src}(\mathrm{self})) \rangle \end{array}} \tag{30}$$

Notes: A modification of the basic add belief rule which also generates an event noting the belief change. Note here that the new intention is given the source "self" not the source of the original intention—this is because any further changes triggered by the belief change are dependent on the agent's internal reasoning and not on the original source of the belief change.

16. HandleAddGoal (Abstract)

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = +!_{\tau_g}g}{\mathrm{undefined}} \tag{31}$$

Notes: Sets up the necessary preconditions for handling an event involving the addition of a goal but doesn't define any transition.

17. HandleAddPerformGoal

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = +!_{\tau_g}g}{\langle ag, i \rangle \rightarrow \langle ag, i' = (+!_{\tau_g}g, \theta^{\mathrm{hd}(i)} \cup \theta_b);_p(+!\mathrm{hd}_e, \mathit{null}, \theta^{\mathrm{hd}(i)});_p i \rangle} \tag{32}$$

Notes: Commit to a perform goal. Unlike `HandleAddAchieveTestGoal` there are no checks for whether we believe the goal. We leave a null action on the stack though so we don't loose track of the previous event (which can happen if that event was planned with only one perform goal).

18. HandleBelief (Abstract)

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = ?b}{\mathrm{undefined}} \tag{33}$$

Notes: Sets up the necessary preconditions for handling a belief (either addition or deletion) but doesn't define the transition. The preconditions are that the top guard on the intention is believed and that the top deed is off belief type. We use ? here to show that the abstract rule does not differentiate between adding and removing a belief.

19. HandleDelayedAction

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a}{\langle ag, i, A \rangle \rightarrow \langle ag, i' = \mathrm{tl}_i(i)\mathrm{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), A' = A; a \rangle} \tag{34}$$

Notes: This rule is intended for use in languages which use an action queue to store actions as they are processed, but delay their actual execution in the environment until some later point. This rule is untested.

20. HandleDropBelief

$$
\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = -b}{B^1 = \{b'|b' \in B \land \mathit{unify}(b', b) \land \mathbf{relevant}(\mathrm{src}(b'), \mathrm{src}(b))\}}
{\langle ag, B, i, I \rangle \to \langle ag, B' = B \backslash B^1, i' = \mathrm{tl}_i(i)\mathrm{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), I \rangle} \tag{35}
$$

Notes: This is the basic rule for dropping beliefs from the belief base. If the deed, $-b$ is on top of the current intention all beliefs that unify with b are removed from the belief base.

21. HandleDropBeliefwEvent

$$
\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = -b \\ B^1 = \{b'|b' \in B \land \mathit{unify}(b', b) \land \mathbf{relevant}(\mathrm{src}(b'), \mathrm{src}(b))\} \end{array}}{\begin{array}{c} \langle ag, B, i, I \rangle \to \\ \mathit{wake}(\langle ag, B' = B \backslash B^1, i' = \mathrm{tl}_i(i)\mathrm{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), I' = I \cup \mathrm{new}(-b, \mathrm{src}(i)) \rangle) \end{array}} \tag{36}
$$

Notes: This rule drops a belief from the belief base (providing the source of the instruction to drop the belief is deemed "relevant" to the source of the belief). At the same time it generates a new intention containing the event that the belief has been dropped. Appropriate handling of this event can allow the agent to form plans in reaction to it.

22. HandleDropGeneralGoal

$$
\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = -!_{\tau_g}g \quad \tau_g \notin \mathit{Exclusions} \\ \exists e \in \mathrm{events}(i).\mathit{unify}(e, +!_{\tau_g}g) \end{array}}{\langle ag, i \rangle \to \langle ag, i' = i.\mathrm{drop}_E(e) \rangle} \tag{37}
$$

$$
\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = -!_{\tau_g}g \quad \tau_g \notin \mathit{Exclusions} \\ \neg\exists e \in \mathrm{events}(i).\mathit{unify}(e, +!_{\tau_g}g) \quad I_1 = \{i'|i' \in I \land \exists e' \in i'.\mathit{unify}(e', +!_{\tau_g}g)\} \end{array}}{\langle ag, I \rangle \to \langle ag, I' = I \cup \{i''|\exists i' \in I_1 \land i'' = i'.\mathrm{drop}_E(e)\} \backslash I_1 \rangle} \tag{38}
$$

Notes: This rule can be parameterised by a set of excluded goal types (*Exclusions*) above. This allows it to act as a default rule for handling "drop goal" instructions. It does this by searching the current intention for the most recent add goal event that unifies with the goal to be dropped and then deletes all rows on the intention above that. drop_E isn't a built-in AIL function but is created by recursing through the intention's events. The second rule drops the event from all the intentions that contain it, assuming the goal doesn't occur in the current intention.

23. HandleDropGoal (Abstract)

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = -!_{\tau_g}g}{\mathrm{undefined}} \quad (39)$$

Notes: Sets up the necessary preconditions for handling a drop goal event but doesn't define the transition.

24. HandleEmptyDeedStack

$$\frac{\mathtt{empty}(i)}{\langle ag, i \rangle \to \langle ag, i \rangle} \quad (40)$$

Notes: Does nothing if the current intention's deed stack is empty.

25. HandleGeneralAction

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \xi.\mathbf{do}(a\theta_b) = \theta_a \quad \tau_a(a) \notin Excluded}{\langle ag, i \rangle \to \langle ag, i' = \mathtt{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_a) \rangle} \quad (41)$$

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \neg\xi.\mathbf{do}(a\theta_b) \quad \mathrm{hd}_e(i) = +!_{\tau_g}g \quad \tau_a(a) \notin Excluded}{\langle ag, i \rangle \to \langle ag, i' = \mathtt{new}(\mathtt{x}!_{\tau_g}g, \theta^{\mathrm{hd}(i)} \cup \theta_a);_p i) \rangle} \quad (42)$$

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \neg\xi.\mathbf{do}(a\theta_b) \quad \neg\mathrm{hd}_e(i) = +!_{\tau_g}g \quad \tau_a(a) \notin Excluded}{\langle ag, i \rangle \to \langle ag, i' = \mathtt{tl}_i(i)\mathtt{U}_\theta\theta^{\mathrm{hd}(i)} \rangle} \quad (43)$$

Notes: HandleGeneralAction extends HandleActionwProblem. The rule is parameterised by a set of action types that are "Excluded"—for instance the GWENDOLEN implementation handles "send" actions differently to other types of action so these are excluded from consideration by this particular rule.

26. HandleGeneralDelayedAction

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = a \quad \tau_a(a) \notin Excluded}{\langle ag, i, A \rangle \to \langle ag, i' = \mathtt{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), A' = A; a \rangle} \quad (44)$$

Notes: This extends HandleDelayedAction in the same way that HandleGeneralAction extends HandleActionwProblem. That is the rule can be parameterised with a list of "Excluded" action types which are to be handled by alternative rules.

27. HandleGoal (Abstract)

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = ?!_{\tau_g}g}{\mathrm{undefined}} \quad (45)$$

Notes: Sets up the necessary preconditions for handling an event involving a goal but doesn't define any transition.

28. HandleGuardNotSatisfied

$$\frac{ag \not\models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}}{A \to A} \tag{46}$$

Notes: The agent does nothing if the guard on an intention can not be satisfied. Should be used with caution since it can cause programs to loop.

29. HandleLockUnlock

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = \mathtt{lock}}{\langle ag, i \rangle \to \langle ag, i' = \mathtt{lock}(\mathtt{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b)) \rangle} \tag{47}$$

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = \mathtt{unlock}}{\langle ag, i \rangle \to \langle ag, i' = \mathtt{unlock}(\mathtt{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b)) \rangle} \tag{48}$$

Notes: This allows an intention to be "locked" as the current intention, for instance to allow a complete sequence of belief changes be processed before any other reasoning takes place. Once finished the intention has to be unlocked. See the SelectIntention rule to see how this works. In languages without the lock construct locking isn't used, obviously, and so doesn't affect intention selection.

30. HandleMessages

$$\frac{}{\begin{array}{c} \langle ag, I, In \rangle \to \\ \langle ag, I' = I \cup \{\mathtt{new}(+received(m), \mathtt{src}(ag)) \mid \downarrow^{ag} m \in In, In' = [] \} \rangle \end{array}} \tag{49}$$

Notes: This rule does not poll the environment for messages. It takes all messages currently in an agent's inbox and converts them to intentions (triggered by a perception that the message has been received), emptying the inbox in the process. It should be noted that it does not store the message anywhere once the inbox is emptied. It assumes that some plan will act appropriately to the message received event. If this does not happen then the message content may be lost.

31. HandleNull

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} \cup \theta_b = null}{\langle ag, i \rangle \to \langle ag, i' = \mathtt{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b) \rangle} \tag{50}$$

Notes: The null action is used as a place holder to preserve, in some situations, a record of an event in an intention stack. This rule simply ignores the null action when it is encountered.

32. HandleSendAction

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = \uparrow^{ag} m \quad \xi.\mathbf{do}((\uparrow^{ag} m)\theta_b) = \theta_a}{\begin{array}{c} \langle ag, i, I, Out \rangle \to \\ \langle ag, i' = \mathtt{tl}_i(i)\mathtt{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_a), I'I \cup \{\mathtt{new}(+\uparrow^{ag} m, \mathtt{src}(\mathtt{self})) \}, \\ Out' = Out \cup \{m\} \rangle \end{array}} \tag{51}$$

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = \uparrow^{ag} m \quad \neg\xi.\mathbf{do}((\uparrow^{ag} m)\theta_b) \quad \mathrm{hd}_e(i) = +!_{\tau_g} g}{\langle ag, i, I, Out\rangle \rightarrow}$$
$$\langle ag, i' = (\mathrm{x}!_{\tau_g} g, \theta^{\mathrm{hd}(i)} \cup \theta_b);_p i, I' = I \cup \{\mathrm{new}(+\uparrow^{ag} m, \mathrm{src}(\mathrm{self}))\},$$
$$Out' = Out \cup \{m\}\rangle$$

$$(52)$$

Notes: This rule is implemented as an extension of HandleActionwProblem. As well as executing a send action it also adds a record the message has been sent to the outbox and generates an intention from the message sending event.

33. HandleTopDeed (Abstract)

$$\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b}{\mathrm{undefined}} \tag{53}$$

Notes: Sets up the necessary preconditions for handling the top of the deed stack but doesn't define the transition.

34. HandleWaitFor

$$\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = *b \\ P = \xi.\mathbf{Percepts}(ag) \quad OP = oldPercepts(P) \quad b \in B \cup P \backslash OP \end{array}}{\begin{array}{c} \langle ag, B, i, I, In\rangle \rightarrow \\ wake(\langle ag, B' = B \cup P \backslash OP, i' = \mathrm{tl}_i(i)\mathrm{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), \\ I' = I \cup \{\mathrm{new}(+b') \mid b' \in P \wedge b \notin B\} \cup \{\mathrm{new}(-b') \mid b' \in OP\}, \\ In' = In \cup \xi.\mathbf{getMessages}()\rangle) \end{array}} \tag{54}$$

$$\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = *b \\ P = \xi.\mathbf{Percepts}(ag) \quad OP = oldPercepts(P) \quad b \notin B \cup P \backslash OP \\ (\neg\mathbf{allsuspended} \vee \mathbf{P} \neq \emptyset) \end{array}}{\begin{array}{c} \langle ag, B, i, I, In\rangle \rightarrow \\ wake(\langle ag, B' = B \cup P \backslash OP, i' = \mathrm{suspend}(i) \\ I' = I \cup \{\mathrm{new}(+b') \mid b' \in P \wedge b \notin B\} \cup \{\mathrm{new}(-b') \mid b' \in OP\} \\ In' = In \cup \xi.\mathbf{getMessages}()\rangle) \end{array}} \tag{55}$$

$$\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = *b \\ b \notin B \quad \mathbf{allsuspended} \quad \xi.\mathbf{Percepts}(ag) = \emptyset \end{array}}{\begin{array}{c} \langle ag, B, In, i\rangle \rightarrow \\ \mathbf{sleep}(\langle ag, B, In \cup \xi.\mathbf{getMessages}(), \mathrm{suspend}(i)\}\rangle) \end{array}} \tag{56}$$

Notes: Because sleeping and waiting behaviour is critical to model checking with JPF we found it necessary to introduce new syntax which allows an intention to wait until the agent holds a certain belief. This syntax is $*b$. If the relevant belief is not held then that intention is suspended and if all intentions are suspended then the agent is told to sleep at the next opportunity. This rule for handling the waiting behaviour is actually implemented as an extension to perception and so makes the decision to suspend the intention based on the most recent information available to the agent. It should be noted that several other

rules, e.g. ones that add new beliefs to the belief base automatically unsuspended all intentions allowing the wait for deed to be rechecked.

35. HandleWaitForDirect

$$
\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = *b \\ P = \xi.\textbf{Percepts}(ag) \quad OP = oldPercepts(P) \quad b \in B \cup P \backslash OP \end{array}}{\begin{array}{c} \langle ag, B, i, In \rangle \to \\ wake(\langle ag, B' = B \cup P \backslash OP, i' = \mathrm{tl}_i(i)\mathrm{U}_\theta(\theta^{\mathrm{hd}(i)} \cup \theta_b), \\ In' = In \cup \xi.\textbf{getMessages}()\rangle) \end{array}} \tag{57}
$$

$$
\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = *b \\ P = \xi.\textbf{Percepts}(ag) \quad OP = oldPercepts(P) \quad b \notin B \cup P \backslash OP \\ (\neg\textbf{allsuspended} \vee \textbf{P} \neq \emptyset) \end{array}}{\begin{array}{c} \langle ag, B, i, In \rangle \to \\ wake(\langle ag, B' = B \cup P \backslash OP, i' = \mathrm{suspend}(i) \\ In' = In \cup \xi.\textbf{getMessages}()\rangle) \end{array}} \tag{58}
$$

$$
\frac{\begin{array}{c} ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad \mathrm{hd}_d(i)\theta^{\mathrm{hd}(i)} = *b \\ b \notin B \quad \textbf{allsuspended} \quad \xi.\textbf{Percepts}(ag) = \emptyset \end{array}}{\begin{array}{c} \langle ag, B, In, i \rangle \to \\ \textbf{sleep}(\langle ag, B, In \cup \xi.\textbf{getMessages}(), \mathrm{suspend}(i)\}\rangle) \end{array}} \tag{59}
$$

Notes: This rule is essentially the same as HandleWaitFor. The only difference is that it sub-classes DirectPerception rather than Perceive. This means it does not create new intentions from incoming belief changes.

36. IgnoreUnplannedProblemGoal

$$
\frac{\mathrm{hd}_e(i) = \mathrm{x}!_{\tau_g} g \quad \mathrm{hd}_d(i) = \varepsilon}{\langle ag, i \rangle \to \langle ag, i \rangle} \tag{60}
$$

Notes: This rule ignores an unplanned problem goal. It simply does nothing but allows the reasoning cycle of the agent to continue processing on the assumption that planning of the goal may become possible later.

37. MatchDropGoal

$$
\frac{\begin{array}{c} \mathrm{deeds}(i) = [\varepsilon] \quad \mathrm{hd}_e = \mathrm{x}!_{\tau_g} g \\ I' = \{\langle i', \theta \rangle \mid i' \in I \wedge \exists +!g' \in \mathrm{events}(i').\, unify(g', g) = \theta \wedge \mathrm{src}(i') = \mathrm{src}(i)\} \\ I'' = \{(-!_{\tau_g}g, \varepsilon, T, \theta);_p i \mid \langle i, \theta \rangle \in I'\} \quad \mathscr{S}_{\mathrm{int}}(I \backslash I' \cup I'') = (i_1, I_1) \end{array}}{\langle ag, i, I \rangle \to \langle ag, i' = i_1, I' = I_1 \rangle} \tag{61}
$$

Notes: If an intention consists only of a single unplanned problem goal event then this should be matched with and placed as a row on the top of all intentions containing that goal. The intention is then removed and a new current intention selection. This is intended to assist in dealing with requests to drop goals that come from some external source.

38. Perceive

$$
\frac{ag \models \mathrm{hd}_g(i)\theta^{\mathrm{hd}(i)}, \theta_b \quad P = \xi.\textbf{Percepts}(ag) \quad OP = oldPercepts(P)}{\begin{array}{c} \langle ag, B, I, In \rangle \rightarrow \\ wake(\langle ag, B' = B \cup P \backslash OP, \\ I \cup \{\mathtt{new}(+b) \mid b \in P) \wedge b \notin B\} \cup \{\mathtt{new}(-b) \mid b \in OP\}, \\ In \cup \xi.\textbf{getMessages}()\rangle) \end{array}} \quad (62)
$$

Notes: This is a complex perception rule. It adds all percepts to the belief base and removes all beliefs no longer perceived. It also add all messages to the inbox. It then creates new intentions, each triggered by the events of acquiring or losing one of the percepts. A key part of the working of the rule depends on AIL's annotation of all beliefs in the belief base with a source and its use of a special annotation for beliefs whose source is perception.

39. ProcessDelayedAction

$$
\frac{\xi.\textbf{do}(a)}{\langle ag, A = a; A' \rangle \rightarrow \langle ag, A' \rangle} \quad (63)
$$

Notes: This rule executes the top action on the agent's action queue. The assumption is that a (FIFO) queue of actions has been created which the agent will execute in order at some point. During planning actions are put on the queue but not actually executed. This rule is untested.

40. SelectIntention

$$
\frac{\neg\mathtt{empty}(i) \quad \neg\mathtt{locked}(i) \quad \mathscr{S}_{\mathrm{int}}(I \cup \{i\}) = (i', I') \quad \neg\textbf{allsuspended}}{\langle ag, i, I \rangle \rightarrow \langle ag, i', I' \rangle} \quad (64)
$$

$$
\frac{\neg\mathtt{empty}(i) \quad \mathtt{locked}(i) \quad \neg\textbf{allsuspended}}{\langle ag, i, I \rangle \rightarrow \langle ag, i, I \rangle} \quad (65)
$$

Notes: This is the basic rule for intention selection. It works by calling the agent method selectIntention ($\mathscr{S}_{\mathrm{int}}$) which is expected to be a common candidate for over-riding. The only situation in which $\mathscr{S}_{\mathrm{int}}$ is not called is if the current intention is locked in which case that intention is selected again.

41. SelectIntentionNotUnplannedProblemGoal

$$
\frac{\begin{array}{c} \neg\mathtt{empty}(i) \quad \neg\mathtt{locked}(i) \quad \mathscr{S}_{\mathrm{int}}(I \cup \{i\}) = (i', I') \\ \mathrm{hd}_e(i') \neq -!_{\tau_g}g \vee \neg\mathtt{noplan}(i') \quad \neg\textbf{allsuspended} \end{array}}{\langle ag, i, I \rangle \rightarrow \langle ag, i', I' \rangle} \quad (66)
$$

$$
\frac{\neg\mathtt{empty}(i) \quad \mathtt{locked}(i) \quad \mathrm{hd}_e(i') \neq -!_{\tau_g}g \vee \neg\mathtt{noplan}(i') \quad \neg\textbf{allsuspended}}{\langle ag, i, I \rangle \rightarrow \langle ag, i, I \rangle} \quad (67)
$$

Notes: This rule extends SelectIntention with one additional condition which is that the top event on the selected intention is not a problem or drop goal event which has no associated plan. We assume that, where this rule is used, another rule (e.g. MatchDropGoal) is employed to handle these intentions.

42. SleepIfEmpty

$$\frac{(i = null \lor \texttt{empty}(i) \lor \textbf{is\_suspended}(i)) \land (I = \emptyset \lor \textbf{allsuspended})}{\langle ag, i, I \rangle \to \textbf{sleep}(\langle ag, i, I \rangle)} \quad (68)$$

Notes: This rule sleeps an agent thread in a controlled fashion if all its intention are empty or suspended. It is quite important to include this rule, or one similar to it, into a language semantics even if one isn't there. The JPF model checker does not assume a fair JAVA scheduling algorithm so unless a multi-agent system forces agents (and so their threads) to sleep it will investigate runs in which one agent executes continuously and none of the others do so. Note that this does not immediately sleep the thread. It sets a flag that the agent wishes to sleep. The agent controller then decides when this should happen.

Appendix B: Sample GWENDOLEN code for the verified scenarios

B.1 Notation

GWENDOLEN uses the AIL's plan mechanisms "off the shelf". It makes no use of prefix matching in plan execution and only has plans whose prefix is ε. Similarly, it is assumed that the guard stacks for plans are all $\top$ except for the very top guard that governs the plan applicability. Plans are therefore a triple of a triggering event (if relevant), a guard, and a body of deeds to be performed. We therefore represent GWENDOLEN's version of the plan

trigger	$+!_a\texttt{clean()}$
prefix	$[\varepsilon]$
guard	$\texttt{dirty(Room)}$
	$\top$
body	$+!_a\texttt{goto(Room)}$
	$+!_a\texttt{vacuum(Room)}$

that appeared in Sect. 3.1.8 as

$$+!_a clean() : dirty(Room) \ \texttt{<-} \ +!_a goto(Room); +!_a vacuum(Room) \quad (69)$$

GWENDOLEN agents also distinguish two sorts of goals. *Achievement goals*, $!_a g$, make statements about beliefs the agent wishes to hold. They remain goals until the agent gains an appropriate belief. *Perform goals*, $!_p g$, simply state a sequence of deeds to be performed and cease to be a goal as soon as that sequence is complete. When an agent takes an action, it executes code specific to that action in the environment. Typically, this code alters the set of propositions that agents are able to perceive. It may also cause messages to be added to an agent's inbox. Agents go

through a specific perception phase when they check their beliefs against the environment's percepts and modify them accordingly. At this point, agents also handle the messages currently in their inbox.

$\uparrow (a, p, m)$ indicates the sending of a message m, with performative, p, to agent a, and $\downarrow (p, m)$ indicates the receipt of a message m with performative p. Since guards may also refer to sent and received messages, the syntax $\uparrow (a, p, m)$ is also used in plan guards, as is $\downarrow (p, m)$. In the following examples, the relevant performatives are **p** (for *perform* the message content) and **b** (for *believe* the message content).

$\mathscr{B}b$ is used in plan guards to indicate the condition that the agent believes b. The lock, unlock, and $*b$ deeds (see Sect. 3.1.7) appear as such in the code.

Throughout this section, we will use ',' to indicate concatenation of deeds on a stack. We follow the Prolog convention of representing variables as starting with upper-case letters, while constants start with lower-case letter. The belief rules used for Prolog-style reasoning in guards are represented in a Prolog style as l1 :- l2, l3, l4;

B.2 Contract net code

Action a makes fact g true in the environment. Similarly, action $a2$ makes the $g2$ true. These facts can then be perceived by all agents.

Code Example 2.1 Contract Net

```
:name: ag1                                                                    1
                                                                              2
: Initial   Beliefs :                                                         3
                                                                              4
¬cando(g)                                                                     5
¬cando(g2)                                                                    6
ag(ag2)                                                                       7
ag(ag3)                                                                       8
my_name(ag1)                                                                  9
                                                                             10
: Initial   Goals:                                                           11
                                                                             12
!_a g                                                                        13
!_a g2                                                                       14
                                                                             15
:Plans:                                                                      16
                                                                             17
+!_a g : {𝓑 cando(g)}  ← a;                                                  18
+!_a g2 : {𝓑 cando(g2)}  ← a2;                                               19
+!_a Gl : {𝓑 ¬cando(Gl)}  ← +!_p cfp(Gl);                                    20
+!_p cfp(G) :                                                                21
   {𝓑 ag(A1), 𝓑 my_name(Name), ¬↑ (A1, p, respond(G, Name)), ¬𝓑 proposal(Ag, G)}  22
   ← ↑(A1, p, respond(G, Name));                                             23
+!_p cfp(G1) : {𝓑 proposal(A, G1)} ← *G1;                                    24
+proposal(A, G4) : {𝓑 ¬cando(G4), ¬𝓑 awarded(G4)} ← ↑(A, b, award(G4)),     25
   +awarded(G4);                                                             26
+↓(p, Ga) : ⊤ ← +!_p Ga;                                                     27
+↓(b, B): ⊤ ← +B;                                                            28
+!_p respond(G2, Name) : {𝓑 cando(G2), 𝓑 my_name(A)} ← ↑(Name, b, proposal(A, G2));  29
+!_p respond(G3, Name) : {𝓑 ¬cando(G3), 𝓑 my_name(A)} ← ↑(Name, b, sorry(A, G3));   30
```

+award(G5) : ⊤ ← +!$_a$ G5 ;   31

+!$_p$ cfp(G6) : ⊤ ← *G6;   32

                                                                                     33

:**name**: ag2   34

                                                                                     35

: **Initial   Beliefs** :   36

                                                                                     37

my_name(ag2)   38

cando(g)   39

¬cando(g2)   40

                                                                                     41

:**Plans**:   42

                                                                                     43

+!$_a$ g : {$\mathscr{B}$ cando(g)} ← a;   44

+!$_a$ g2 : {$\mathscr{B}$ cando(g2)} ← a2;   45

+!$_a$ Gl : {$\mathscr{B}$ ¬cando(Gl)} ← +!$_p$ cfp(Gl);   46

+!$_p$ cfp(G) :   47

    {$\mathscr{B}$ ag(A1), $\mathscr{B}$ my_name(Name), ¬↑ (A1, **p**, respond(G, Name)), ¬$\mathscr{B}$ proposal(Ag, G)}   48

    ← ↑(A1, **p**, respond(G, Name));   49

+!$_p$ cfp(G1) : {$\mathscr{B}$ proposal(A, G1)} ← *G1;   50

+proposal(A, G4) : {$\mathscr{B}$ ¬cando(G4), ¬$\mathscr{B}$ awarded(G4)} ← ↑(A, **b**, award(G4)),   51

    +awarded(G4);   52

+↓(**p**, Ga) : ⊤ ← +!$_p$ Ga;   53

+↓(**b**, B): ⊤ ← +B;   54

+!$_p$ respond(G2, Name) : {$\mathscr{B}$ cando(G2), $\mathscr{B}$ my_name(A)} ← ↑(Name, **b**, proposal(A, G2));   55

+!$_p$ respond(G3, Name) : {$\mathscr{B}$ ¬cando(G3), $\mathscr{B}$ my_name(A)} ← ↑(Name, **b**, sorry(A, G3));   56

+award(G5) : ⊤ ← +!$_a$ G5 ;   57

+!$_p$ cfp(G6) : ⊤ ← *G6;   58

                                                                                     59

:**name**: ag3   60

                                                                                     61

: **Initial   Beliefs** :   62

                                                                                     63

my_name(ag3)   64

cando(g2)   65

¬cando(g)   66

                                                                                     67

:**Plans**:   68

                                                                                     69

+!$_a$ g : {$\mathscr{B}$ cando(g)} ← a;   70

+!$_a$ g2 : {$\mathscr{B}$ cando(g2)} ← a2;   71

+!$_a$ Gl : {$\mathscr{B}$ ¬cando(Gl)} ← +!$_p$ cfp(Gl);   72

+!$_p$ cfp(G) :   73

    {$\mathscr{B}$ ag(A1), $\mathscr{B}$ my_name(Name), ¬↑ (A1, **p**, respond(G, Name)), ¬$\mathscr{B}$ proposal(Ag, G)} ←   74

    ↑(A1, **p**, respond(G, Name));   75

+!$_p$ cfp(G1) : {$\mathscr{B}$ proposal(A, G1)} ← *G1;   76

+proposal(A, G4) : {$\mathscr{B}$ ¬cando(G4), ¬ $\mathscr{B}$ awarded(G4)} ← ↑(A, **b**, award(G4)),   77

    +awarded(G4);   78

+↓(**p**, Ga) : ⊤ ← +!$_p$ Ga;   79

+↓(**b**, B): ⊤ ← +B;   80

+!$_p$ respond(G2, Name) : {$\mathscr{B}$ cando(G2), $\mathscr{B}$ my_name(A)} ← ↑(Name, **b**, proposal(A, G2));   81

+!$_p$ respond(G3, Name) : {$\mathscr{B}$ ¬cando(G3), $\mathscr{B}$ my_name(A)} ← ↑(Name, **b**, sorry(A, G3));   82

+award(G5) : ⊤ ← +!$_a$ G5 ;   83

+!$_p$ cfp(G6) : ⊤ ← *G6;   84

B.3 Basic auction code

Code Example 2.2 Basic Auction

```
:name: ag1                                                                    1
                                                                              2
: Initial  Beliefs :                                                          3
                                                                              4
:Plans:                                                                       5
                                                                              6
+↓(b, B):  ⊤  ← +B;                                                           7
+bid(Z, A) : {ℬ bid(X1, ag2), ℬ bid(X2, ag3), ℬ bid(X3, ag4), ℬ bid(200, Ag)} ←   8
  ↑(Ag, b, win);                                                              9
                                                                              10
:name: ag2                                                                    11
                                                                              12
: Initial  Beliefs :                                                          13
                                                                              14
my_name(ag2)                                                                  15
                                                                              16
: Initial  Goals:                                                             17
                                                                              18
!_p bid                                                                       19
                                                                              20
:Plans:                                                                       21
                                                                              22
+↓(b, B):  ⊤  ← +B;                                                           23
+!_p bid : {ℬ my_name(Name), ¬↑(ag1, b, bid(100, Name))} ← ↑(ag1, b, bid(100, Name));   24
```

Agents 3 and 4 are identical to agent 2 except for the amount they bid.

B.4 Auction coalition code

The code for agent 1 is the same as in the Basic Auction example.

Code Example 2.3 Auction Coaltion

```
:name: ag2                                                                    1
                                                                              2
: Initial  Beliefs :                                                          3
                                                                              4
my_name(ag2)                                                                  5
                                                                              6
: Initial  Goals:                                                             7
                                                                              8
!_p coalition                                                                 9
                                                                              10
:Plans:                                                                       11
                                                                              12
+↓(b, B):  ⊤  ← +B;                                                           13
+!_p bid : {ℬ my_name(Name), ¬↑(ag1, b, bid(250, Name))} ← ↑(ag1, b, bid(250, Name));   14
+!_p coalition : {ℬ my_name(Ag), ¬↑(ag3, b, coalition (Ag))} ← ↑(ag3, b, coalition (Ag));   15
+agree(A, X) :  ⊤  ← +!_p bid;                                                16
                                                                              17
:name: ag3                                                                    18
                                                                              19
```

```
: Initial  Beliefs :                                                          20
                                                                              21
my_name(ag3)                                                                  22
                                                                              23
: Initial  Goals:                                                             24
                                                                              25
!_p bid                                                                       26
                                                                              27
: Plans:                                                                      28
                                                                              29
+↓(b, B):  ⊤  ← +B;                                                           30
+!_p _p bid : {𝓑 my_name(Name), ¬↑(ag1, b, bid(150, Name))} ←                31
                 ↑(ag1, b, bid(150, Name));                                   32
+ coalition (A) : {𝓑 my_name(Name), ¬↑(A, b, agree(Name, 150))} ←            33
                 ↑(A, b, agree(Name, 150));                                   34
```

The code for agent 4 is the same as for agent 3 except with a different bid amount.

B.5 Dynamic auction coalition code

The action *win*(*Z*, *A*) makes the fact that agent *A* has won with amount *Z* available to all agents by perception.

Code Example 2.4 Dynamic Auction

```
: name: ag1                                                                   1
                                                                              2
: Initial  Beliefs :                                                          3
                                                                              4
my_name(ag1)                                                                  5
                                                                              6
: Belief  Rules:                                                             7
                                                                              8
𝓑  allbids  :− 𝓑 bid_processed(ag2), 𝓑 bid_processed(ag3), 𝓑 bid_processed(ag4);    9
                                                                              10
: Plans:                                                                      11
                                                                              12
+↓(b, bid(D, From)) : {𝓑 bid(E, From)} ← −bid(From, E),                      13
    +bid(From, D);                                                            14
+↓(b, bid(D, From)) : {¬𝓑 bid(E, From)} ← +bid(From, D);                     15
+bid(Z, A) : {𝓑 current_winner(Ag1, Amw), Amw < A, 𝓑allbids} ← lock,          16
    −current_winner(Ag1, Amw),                                                17
    +ann_winner,                                                              18
    +current_winner(Z, A),                                                    19
    win(Z, A),                                                                20
    unlock;                                                                   21
+bid_processed(Ag) : {𝓑 current_winner(Agw, Amw), 𝓑allbids , ¬ 𝓑  ann_winner} ←    22
    lock,                                                                     23
    +ann_winner,                                                             24
    win(Agw, Amw),                                                            25
    unlock;                                                                   26
+bid(Ag, Am) : {¬𝓑 current_winner(Ag2, Amw)} ←                               27
    +current_winner(Ag, Am),                                                  28
    +bid_processed(Ag);                                                       29
+bid(Ag, Am) : {𝓑 current_winner(Agw, Amw), ¬(Am < Amw), ¬𝓑allbids } ← lock,   30
    +current_winner(Ag, Am),                                                  31
    +bid_processed(Ag),                                                       32
```

```
      −current_winner(Agw, Amw),                                              33
      unlock;                                                                 34
+bid(Ag, Am) : {ℬ current_winner (Agw, Amw), Am < Amw, ¬ℬallbids} ←          35
      +bid_processed(Ag);                                                     36
                                                                             37
:name: ag2                                                                   38
                                                                             39
: Initial   Beliefs :                                                        40
                                                                             41
my_name(ag2)                                                                 42
 collaborator (ag3)                                                          43
cash(150)                                                                    44
                                                                             45
: Initial   Goals:                                                           46
                                                                             47
!ₚ bid                                                                       48
                                                                             49
:Plans:                                                                      50
                                                                             51
+↓(b, B): ⊤ ← +B;                                                            52
+!ₚ bid : {ℬ my_name(Name), ℬcash(C), ¬ ↑(ag1, b, bid(C, Name))} ←          53
      ↑(ag1, b, bid(C, Name));                                               54
+agree(A, X): {ℬ cash(C), ℬ my_name(Name)} ← ↑(ag1, b, bid((C + X), Name));  55
+win(Ag, X): {ℬ my_name(Name), ¬ℬwin(Name, Any), ℬ collaborator (Coll)} ←   56
      +!ₐ coalition (Coll)  ;                                                57
+!ₐ _p coalition (Coll)   : {ℬ my_name(Ag), ¬↑(Coll, b,  coalition (Ag))} ← 58
      ↑(Coll, b,  coalition (Ag)),                                           59
      + coalition (Coll );                                                   60
                                                                             61
:name: ag3                                                                   62
                                                                             63
: Initial   Beliefs :                                                        64
                                                                             65
my_name(ag3)                                                                 66
cash(150)                                                                    67
                                                                             68
: Initial   Goals:                                                           69
                                                                             70
!ₚ bid                                                                       71
                                                                             72
:Plans:                                                                      73
                                                                             74
+↓(b, B): ⊤ ← +B;                                                            75
+!ₚ bid : {ℬ my_name(Name), ℬcash(C), ¬ ↑(ag1, b, bid(C, Name))} ←          76
      ↑(ag1, b, bid(C, Name));                                               77
+ coalition (A) : {ℬ my_name(Name), ℬcash(C), ¬ ↑(A, b, agree(Name, C))} ←  78
      ↑(A, b, agree(Name, C));                                              79
```

The code for agent 4 is the same as for agent 3 except with a different bid amount.

B.6 Auction trust code

The action *win*(A, Z) makes the fact that agent A has won with amount Z available to all agents by perception.

Code Example 2.5 Trust Auction

```
:name: ag1                                                              1
                                                                        2
: Initial   Beliefs :                                                   3
                                                                        4
:Plans:                                                                 5
                                                                        6
+↓(b, bid(D, From)) : {ℬ bid(From, E)} ← −bid(From, E),                 7
   +bid(From, D);                                                       8
+↓(b, bid(D, From)) : {¬ ℬ bid(From, E)} ← +bid(From, D);               9
+bid(Z, A) : {ℬ bid(ag2, X1), ℬ bid(ag3, X2), ℬ bid(ag4, X3),          10
            ¬ ℬ winning_amount(Am), X2 < X1, X3 < X1} ←                 11
   +winning_amount(X1),                                                 12
   win(ag2, X1);                                                        13
+bid(Z, A) :  {ℬ bid(ag2, X1), ℬ bid(ag3, X2), ℬ bid(ag4, X3),         14
            ¬ ℬ winning_amount(Am), X1 < X2, X3 < X2} ←                 15
   +winning_amount(X2),                                                 16
   win(ag3, X2);                                                        17
+bid(Z, A) :  {ℬ bid(ag2, X1), ℬ bid(ag3, X2), ℬ bid(ag4, X3),         18
            ¬ ℬ winning_amount(Am), X2 < X3, X1 < X3} ←                 19
   +winning_amount(X2),                                                 20
   win(ag4, X3);                                                        21
+bid(Z, A) : {ℬ winning_amount(Am), Am < A} ← −winning_amount(Am),      22
   +winning_amount(A),                                                  23
   win(A, Z);                                                           24
                                                                        25
:name: ag2                                                              26
                                                                        27
: Initial   Beliefs :                                                   28
                                                                        29
my_name(ag2)                                                            30
 trust (ag3)                                                            31
                                                                        32
: Initial   Goals:                                                      33
                                                                        34
!_p bid                                                                 35
                                                                        36
:Plans:                                                                 37
                                                                        38
+↓(b, B): ⊤ ← +B;                                                       39
+!_p _p bid : {ℬ my_name(Name), ¬↑(ag1, b, bid(150, Name))} ←           40
   ↑(ag1, b, bid(150, Name));                                           41
+win(A, X) : {ℬ my_name(Name), ¬ℬwin(Name, Y), ℬ trust (Ag),           42
   ¬  ↑(Ag, b,  coalition (Name))} ← ↑(Ag, b,  coalition (Name));      43
+agree(A, X) : ⊤ ← ↑(ag1, b, bid(300, ag2 ));                           44
                                                                        45
:name: ag3                                                              46
                                                                        47
: Initial   Beliefs :                                                   48
                                                                        49
my_name(ag3)                                                            50
                                                                        51
: Initial   Goals:                                                      52
                                                                        53
!_p bid                                                                 54
                                                                        55
                                                                        56
```

```
: Plans:                                                                    57
                                                                            58
+↓(b, B): ⊤ ← +B;                                                           59
+!ₚ bid : {ℬ my_name(Name), ¬↑(ag1, b, bid(150, Name))} ← ↑(ag1, b, bid(150, Name));   60
+ coalition (A) : {ℬ my_name(Name), ¬↑(A, b, agree(Name, 150))} ←           61
   ↑(A, b, agree(Name, 150));                                               62
```

The code for agent 3 is the same as for agent 4 except with a different bid amount.

B.7 Auction dynamic trust code

Code Example 2.6 Dynamic Trust Action

```
: name: ag1                                                                 1
                                                                            2
: Plans:                                                                    3
                                                                            4
+↓(b, bid(D, From)) : {ℬ bid(From, E)} ← −bid(From, E),                     5
   +multiple_bidder (From),                                                 6
   +bid(From, D);                                                           7
+↓(b, bid(D, From)) : {¬ ℬ bid(From, E)} ← +bid(From, D);                   8
+bid(Z, A) : {ℬ bid(ag2, X1), ℬ bid(ag3, X2), ℬ bid(ag4, X3), ℬ bid(ag5, X4),   9
         ¬ ℬ winning_amount(Am), X2 < X1, X3 < X1, X4 < X1} ←               10
   +winning_amount(X1),                                                     11
   win(ag2, X1);                                                           12
+bid(Z, A) : {ℬ bid(ag2, X1), ℬ bid(ag3, X2), ℬ bid(ag4, X3), ℬ bid(ag5, X4),   13
         ¬ ℬ winning_amount(Am), X1 < X2, X3 < X2, X4 < X2} ←               14
   +winning_amount(X2),                                                     15
   win(ag3, X2);                                                           16
+bid(Z, A) : {ℬ bid(ag2, X1), ℬ bid(ag3, X2), ℬ bid(ag4, X3), ℬ bid(ag5, X4),   17
         ¬ ℬ winning_amount(Am), X2 < X3, X1 < X3, X4 < X3} ←               18
   +winning_amount(X3),                                                     19
   win(ag4, X3);                                                           20
+bid(Z, A) : {ℬ bid(ag2, X1), ℬ bid(ag3, X2), ℬ bid(ag4, X3), ℬ bid(ag5, X4),   21
         ¬ ℬ winning_amount(Am), X2 < X4, X1 < X4, X3 < X4} ←               22
   +winning_amount(X4),                                                     23
   win(ag5, X4);                                                           24
+bid(Z, A) : {ℬ winning_amount(Am), Am < A} ← −winning_amount(Am),          25
   +winning_amount(A), win(A, Z);                                           26
+bid(Z, A) : {ℬ multiple_bidder (Z), ℬ winning_amount(Am), A < Am} ←        27
   ↑(Z, b, failed_bid );                                                    28
                                                                            29
: name: ag2                                                                 30
                                                                            31
: Initial   Beliefs :                                                       32
                                                                            33
my_name(ag2)                                                                34
 trust (ag3)                                                                35
 trust (ag4)                                                                36
                                                                            37
: Initial   Goals:                                                          38
                                                                            39
!ₚ bid                                                                      40
                                                                            41
: Plans:                                                                    42
                                                                            43
```

```
+↓(b, B): ⊤ ← +B;                                                                    44
+!ₚ bid : {ℬ my_name(Name), ¬↑(ag1, b, bid(150, Name))} ← ↑(ag1, b, bid(150, Name)); 45
+win(A, X) : {ℬ my_name(Name), ¬ℬwin(Name, Y), ℬ trust (Ag),                          46
        ¬ ℬ formed_coalition (AgB), ¬ ↑(Ag, b, coalition (Name))} ←                   47
    ↑(Ag, b, coalition (Name)),                                                       48
    +formed_coalition (Ag);                                                           49
+failed_bid  : {ℬ my_name(Name), ¬ℬwin(Name, Y),                                      50
        ℬ trust (Ag), ℬ formed_coalition (AgB), ¬ ↑(Ag, b, coalition (Name))} ←       51
    ↑(Ag, b, coalition (Name)),                                                       52
    +formed_coalition (Ag),                                                           53
    −trust (AgB);                                                                     54
+agree(A, X) : ⊤ ← ↑(ag1, b, bid((X + 150), ag2 ));                                   55
                                                                                      56
:name: ag3                                                                            57
                                                                                      58
: Initial  Beliefs :                                                                  59
                                                                                      60
my_name(ag3)                                                                          61
                                                                                      62
: Initial  Goals:                                                                     63
                                                                                      64
!ₚ bid                                                                                65
                                                                                      66
:Plans:                                                                               67
                                                                                      68
+↓(b, B): ⊤ ← +B;                                                                     69
+!ₚ bid : {ℬ my_name(Name), ¬↑(ag1, b, bid(25, Name))} ← ↑(ag1, b, bid(25, Name));    70
+ coalition (A) : {ℬ my_name(Name), ¬↑(A, b, agree(Name, 25))} ←                      71
    ↑(A, b, agree(Name, 25));                                                         72
```

The code for agents 4 and 5 are the same as for agent 3 except with a different bid amount.

B.8 Cleaning robot code

The action *next*(*slot*) moves the robot to the next space on the grid and updates all perceivable facts accordingly (e.g., whether the agent can see any garbage etc.). The action *drop*(*garb*) makes garbage perceivable in that grid slot (assuming the agent was holding garbage). The action *moveTowards*($X1, Y1$) makes the agent move one square towards the coordinates ($X1, Y1$). The action *burn*(*garb*) destroys a piece of garbage.

Code Example 2.7 Cleaning Robots

```
:name: r1                                                                             1
                                                                                      2
: Initial  Beliefs :                                                                  3
                                                                                      4
pos(r2, 2, 2)                                                                         5
checking( slots )                                                                     6
                                                                                      7
:Plans:                                                                               8
                                                                                      9
+pos(r1, X1, Y1) : {ℬ checking( slots ), ¬ ℬ garbage(r1)} ← next( slot );             10
+garbage(r1) : {ℬ checking( slots )} ← +!ₚ stop(check),                               11
```

```
    +!p take(garb, r2),                                             12
    +!p continue(check);                                            13
+!p stop(check) : ⊤ ← +!a pos(r1, X1, Y1),                          14
    +pos(back, X1, Y1),                                             15
    −checking(slots);                                               16
+!p take(S, L) : ⊤ ← +!p ensure_pick(S),                           17
    +!p go(L),                                                      18
    drop(S);                                                        19
+!p ensure_pick(S) : {ℬ garbage(r1)} ← pick(garb),                 20
    +!p ensure_pick(S);                                             21
+!p ensure_pick(S) : ⊤ ← donothing;                                22
+!p continue(check) : ⊤ ← +!p go(back),                            23
    −pos(back, X1, Y1),                                             24
    +checking(slots),                                               25
    next(slot);                                                     26
+!p go(L) : {ℬ pos(L, X1, Y1), ℬ pos(r1, X1, Y1)} ← donothing;     27
+!p go(L) : ⊤ ← +!a pos(L, X1, Y1),                                28
    moveTowards(X1, Y1),                                           29
    +!p go(L);                                                      30
                                                                   31
:name: r2                                                          32
                                                                   33
:Plans:                                                            34
                                                                   35
+garbage(r2) : ⊤ ← burn(garb);                                     36
```

References

de Boer, F.S., Hindriks, K.V., van der Hoek, W., Meyer, J.J.C.: A verification framework for agent programming with declarative goals. J. Appl. Log. **5**(2), 277–302 (2007)

Bond, A.H., Gasser, L. (eds.): Readings in Distributed Artificial Intelligence. Kaufmann, Los Altos (1988)

Bordini, R.H., Visser, W., Fisher, M., Wooldridge, M.: Verifiable multi-agent programs. In: First International Workshop on Programming Multiagent Systems: Languages, Frameworks, Techniques and Tools (ProMAS-03). Lecture Notes in Artificial Intelligence, vol. 3067. Springer, Berlin (2003)

Bordini, R.H., Fisher, M., Visser, W., Wooldridge, M.: Model checking rational agents. IEEE Intell. Syst. **19**(5), 46–52 (2004)

Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.): Multi-Agent Programming: Languages, Platforms and Applications. Springer, Berlin (2005a)

Bordini, R.H., Hübner, J.F., Vieira, R.: Jason and the golden fleece of agent-oriented programming. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) Multi-Agent Programming: Languages, Platforms and Applications, pp. 3–37. Springer, Berlin (2005b). Chap. 1

Bordini, R.H., Fisher, M., Visser, W., Wooldridge, M.: Verifying multi-agent programs by model checking. J. Auton. Agents Multi-Agent Syst. **12**(2), 239–256 (2006)

Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming multi-agent systems in AgentSpeak using Jason. Wiley Series in Agent Technology. Wiley, New York (2007)

Bordini, R.H., Dennis, L.A., Farwer, B., Fisher, M.: Automated verification of multi-agent programs. In: Proc. 23rd IEEE/ACM International Conference on Automated Software Engineering (ASE), pp. 69–78 (2008)

Bordini, R.H., Fisher, M., Wooldridge, M., Visser, W.: Property-based slicing for an agent-oriented programming language. J. Logic Comput. (2009)

Boutilier, C., Shoham, Y., Wellman, M.P.: Economic principles of multi-agent systems. Artif. Intell. **94**(1–2), 1–6 (1997)

Bratman, M.E.: Intentions, Plans, and Practical Reason. Harvard University Press, Cambridge (1987)

Bratman, M.E., Israel, D.J., Pollack, M.E.: Plans and resource-bounded practical reasoning. Comput. Intell. **4**, 349–355 (1988)

Bredin, J., Kotz, D., Rus, D., Maheswaran, R.T., Imer, Ç., Basar, T.: Computational markets to regulate mobile-agent systems. J. Auton. Agents Multi-Agent Syst. **6**(3), 235–263 (2003)

Burch, J.R., Clarke, E.M., McMillan, K.L., Dill, D.L., Hwang, L.J.: Symbolic model checking: 10^{20} states and beyond. Inf. Comput. **98**(2), 142–170 (1992)

Cimatti, A., Clarke, E.M., Guinchiglia, E., Guinchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV 2: An OpenSource tool for symbolic model checking. In: Proc. International Conference on Computer-Aided Verification (CAV 2002). Lecture Notes in Computer Science. Springer, Berlin (2002)

Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press, Cambridge (1999)

Cohen, P.R., Levesque, H.J.: Intention is choice with commitment. Artif. Intell. **42**, 213–261 (1990)

Collins, J., Faratin, P., Parsons, S., Rodríguez-Aguilar, J.A., Sadeh, N.M., Shehory, O., Sklar, E. (eds.): Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis (AMEC/TADA). Lecture Notes in Business Information Processing, vol. 13. Springer, Berlin (2009)

Corera, J.M., Laresgoiti, I., Jennings, N.R.: Using archon. Part 2. Electricity transportation management. IEEE Intell. Syst. **11**(6), 71–79 (1996)

D'Amorim, M.: Efficient explicit-state model checking for programs with dynamically allocated data. Ph.D. thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA (2007). Adviser-Marinov, Darko

Dash, R.K., Vytelingum, P., Rogers, A., David, E., Jennings, N.R.: Market-based task allocation mechanisms for limited-capacity suppliers. IEEE Trans. Syst. Man Cybern., Part A, Syst. Hum. **37**(3), 391–405 (2007)

Dastani, M., van Riemsdijk, M.B., Meyer, J.J.C.: Programming multi-agent systems in 3APL. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) Multi-Agent Programming: Languages, Platforms and Applications, pp. 39–67. Springer, Berlin (2005). Chap. 2

Dastani, M., Tinnemeier, N.A.M., Meyer, J.J.C.: A programming language for normative multi-agent systems. In: Dignum, V. (ed.) Multi-Agent Systems: Semantics and Dynamics of Organizational Models, chap 16. IGI Global (2009)

Davis, R., Smith, R.G.: Negotiation as a metaphor for distributed problem solving. Artif. Intell. **20**(1), 63–109 (1983)

Dennis, L.A., Farwer, B.: Gwendolen: a BDI language for verifiable agents. In: Löwe, B. (ed.) Logic and the Simulation of Interaction and Reasoning AISB'08 Workshop. AISB, Aberdeen (2008)

Dennis, L.A., Fisher, M.: Programming verifiable heterogeneous agent systems. In: Proc. 6th International Workshop on Programming in Multi-Agent Systems (ProMAS). Lecture Notes in Computer Science, vol. 5442, pp. 40–55. Springer, Berlin (2008)

Dennis, L.A., Farwer, B., Bordini, R.H., Fisher, M., Wooldridge, M.: A common semantic basis for BDI languages. In: Proc. 7th International Workshop on Programming Multiagent Systems (ProMAS). Lecture Notes in Artificial Intelligence, vol. 4908, pp. 124–139. Springer, Berlin (2008a)

Dennis, L.A., Hepple, A., Fisher, M.: Language constructs for multi-agent programming. In: Proc. 8th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA). Lecture Notes in Artificial Intelligence, vol. 5056, pp. 137–156. Springer, Berlin (2008b)

Doghri, I.: Formal verification of WAHS: an autonomous and wireless P2P auction handling system. In: Proc. 8th International Conference on New Technologies in Distributed Systems (NOTERE), pp. 1–10. ACM, New York (2008)

Durfee, E.H., Lesser, V.R., Corkill, D.D.: Trends in cooperative distributed problem solving. IEEE Trans. Knowl. Data Eng. **1**(1), 63–83 (1989)

Emerson, E.A.: Temporal and modal logic. In: van Leeuwen, J. (ed.) Handbook of Theoretical Computer Science, pp. 996–1072. Elsevier, Amsterdam (1990)

Fisher, M., Hepple, A.: Executing logical agent specifications. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) Multi-Agent Programming: Languages, Tools and Applications, pp. 1–27. Springer, Berlin (2009)

Fisher, M., Ghidini, C.: Executable specifications of resource-bounded agents. J. Auton. Agents Multi-Agent Syst. **21**(3), 368–396 (2010)

Fisher, M., Singh, M.P., Spears, D.F., Wooldridge, M.: Logic-based agent verification (Editorial). J. Appl. Log. **5**(2), 193–195 (2007)

Fortnow, L., Riedl, J., Sandholm, T. (eds.): Proc. 9th ACM Conference on Electronic Commerce (EC). ACM, New York (2008)

Gabbay, D., Kurucz, A., Wolter, F., Zakharyaschev, M.: Many-dimensional modal logics: theory and applications. Studies in Logic and the Foundations of Mathematics, vol. 148. Elsevier, Amsterdam (2003)

Gibney, M.A., Jennings, N.R., Vriend, N.J., Griffiths, J.M.: Market-based call routing in telecommunications networks using adaptive pricing and real bidding. In: Proc. 3rd International Workshop on Intelligent Agents for Telecommunication Applications (IATA). Lecture Notes in Computer Science, vol. 1699, pp. 46–61. Springer, Berlin (1999)

Haque, N., Jennings, N.R., Moreau, L.: Resource allocation in communication networks using market-based agents. Knowl.-Based Syst. **18**(4–5), 163–170 (2005)

Hierons, R.M., Bogdanov, K., Bowen, J.P., Cleaveland, R., Derrick, J., Dick, J., Gheorghe, M., Harman, M., Kapoor, K., Krause, P., Lüttgen, G., Simons, A.J.H., Vilkomir, S.A., Woodward, M.R., Zedan, H.: Using formal specifications to support testing. ACM Comput. Surv. **41**(2), 1–76 (2009)

Himoff, J., Skobelev, P., Wooldridge, M.: MAGENTA Technology: Multi-agent systems for industrial logistics. In: Proc. 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 60–66. ACM Press, New York (2005)

Hindriks, K., de Boer, F., van der Hoek, W., Meyer, J.J.: Formal semantics for an abstract agent programming language. In: Intelligent Agents IV: Proc. 4th International Workshop on Agent Theories, Architectures and Languages. Lecture Notes in Artificial Intelligence, vol. 1365, pp. 215–229. Springer, Berlin (1998)

Hindriks, K., de Boer, F., van der Hoek, W., Meyer, J.J.: Agent programming with declarative goals. In: Intelligent Agents VII: Proc. 6th International Workshop on Agent Theories, Architectures, and Languages (ATAL). Lecture Notes in Artificial Intelligence, vol. 1986, pp. 228–243. Springer, Berlin (2001)

Hirsch, B., Fricke, S., Kroll-Peters, O., Konnerth, T.: Agent programming in practise—experiences with the JIAC IV agent framework. In: Proc. Workshop "From Agent Theory to Agent Implementation" (AT2AI) (2008)

Holzmann, G.J.: The Spin Model Checker: Primer and Reference Manual. Addison-Wesley, Reading (2003)

Jenkins, B.: Hash Functions. Dr. Dobbs Journal, September (1997)

Jennings, N.R., Wooldridge, M.: Applications of agent technology. In: Agent Technology: Foundations, Applications, and Markets. Springer, Heidelberg (1998)

Jongmans, S.S., Hindriks, K., van Riemsdijk, M.: Model checking agent programs by using the program interpreter. In: Dix, J., Leite, J.a., Governatori, G., Jamroga, W. (eds.) Computational Logic in Multi-Agent Systems. Lecture Notes in Computer Science, vol. 6245, pp. 219–237. Springer, Berlin (2010). doi:10.1007/978-3-642-14977-1_17

Java PathFinder: (2009). http://javapathfinder.sourceforge.net

Kacprzak, M., Lomuscio, A., Penczek, W.: Verification of multiagent systems via unbounded model checking. In: Proc. 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 638–645. IEEE Computer Society, Los Alamitos (2004)

Klemperer, P.: Auctions: Theory and Practice. Princeton University Press, Princeton (2004). See also http://www.nuff.ox.ac.uk/users/klemperer/VirtualBook/VBCrevisedv2.asp

Klügl, F., Bazzan, A., Ossowski, S. (eds.): Applications of Agent Technology in Traffic and Transportation. Whitestein Series in Software Agent Technologies and Autonomic Computing. Birkhäuser, Basel (2005)

Konishi, H., Ray, D.: Coalition formation as a dynamic process. J. Econ. Theory **110**(1), 1–41 (2003)

Ljunberg, M., Lucas, A.: The OASIS air traffic management system. In: Proc. 2nd Pacific Rim International Conference on AI (PRICAI) (1992)

Moreno, A., Garbay, C.: Software agents in health care. Artif. Intell. Med. **27**(3), 229–232 (2003)

Muscettola, N., Nayak, P.P., Pell, B., Williams, B.: Remote agent: to boldly go where no AI system has gone before. Artif. Intell. **103**(1–2), 5–48 (1998)

Owre, S., Shankar, N.: A brief overview of pvs. In: Proc. 21st International Conference on Theorem Proving in Higher Order Logics (TPHOLs). Lecture Notes in Computer Science, vol. 5170, pp. 22–27. Springer, Berlin (2008)

Pokahr, A., Braubach, L., Lamersdorf, W.: Jadex: A BDI Reasoning Engine. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) Multi-Agent Programming: Languages, Platforms and Applications, pp. 149–174. Springer, Berlin (2005)

Raimondi, F., Lomuscio, A.: Automatic verification of multi-agent systems by model checking via ordered binary decision diagrams. J. Appl. Log. **5**(2), 235–251 (2007)

Rao, A.: AgentSpeak(L): BDI agents speak out in a logical computable language. In: Proc. 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW). Lecture Notes in Computer Science, vol. 1038, pp. 42–55. Springer, Berlin (1996)

Rao, A.S., Georgeff, M.P.: An abstract architecture for rational agents. In: Rich, C., Swartout, W., Nebel, B. (eds.) Proceedings of Knowledge Representation and Reasoning (KR&R-92), pp. 439–449 (1992)

Rao, A.S., Georgeff, M.: BDI agents: from theory to practice. In: Proc. 1st International Conference on Multi-Agent Systems (ICMAS), San Francisco, USA, pp. 312–319 (1995)

Reeves, D.M., Wellman, M.P., MacKie-Mason, J.K., Osepayshvili, A.: Exploring bidding strategies for market-based scheduling. Decis. Support Syst. **39**(1), 67–85 (2005)

Rimassa, G., Burmeister, B.: Achieving business process agility in engineering change management with agent technology. In: Proc. 8th AI*IA/TABOO Joint Workshop "From Objects to Agents"—Agents and Industry: Technological Applications of Software Agents (WOA), pp. 1–7. Seneca Edizioni, Torino (2007)

Sandholm, T., Lesser, V.R.: Coalitions among computationally bounded agents. Artif. Intell. **94**(1–2), 99–137 (1997)

Shoham, Y.: Agent-oriented programming. Artif. Intell. **60**(1), 51–92 (1993)

Sistla, A.P., Vardi, M., Wolper, P.: The complementation problem for Büchi automata with applications to temporal logic. Theor. Comput. Sci. **49**, 217–237 (1987)

Smith, R.G.: A Framework for Distributed Problem Solving. UMI Research Press, Ann Arbor (1980)

Smith, R.G., Davis, R.: Frameworks for cooperation in distributed problem solving. IEEE Trans. Syst. Man Cybern. **11**(1), 61–70 (1980)

van Riemsdijk, B., van der Hoek, W., Meyer, J.J.: Agent programming in dribble: from beliefs to goals with plans. In: Proc. 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 393–400. ACM, New York (2003)

Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. Eur. J. Finance **16**(1), 8–37 (1961)

Visser, W., Havelund, K., Brat, G.P., Park, S., Lerda, F.: Model checking programs. Autom. Softw. Eng. **10**(2), 203–232 (2003)

Walsh, W.E., Wellman, M.P.: A market protocol for decentralized task allocation. In: Proc. 3rd International Conference on Multiagent Systems (ICMAS), pp. 325–332. IEEE Computer Society, Los Alamitos (1998)

Webster, M.P., Dennis, L.A., Fisher, M.: Model-checking auctions, coalitions and trust. Tech. Rep. ULCS-09-004, Department of Computer Science, University of Liverpool (2009). http://www.csc.liv.ac.uk/research

Winikoff, M.: JACK™ intelligent agents: an industrial strength platform. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) Multi-Agent Programming: Languages, Platforms and Applications, pp. 175–193. Springer, Berlin (2005). Chap. 7

Winikoff, M.: Implementing commitment-based interactions. In: Proc. 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 1–8. ACM, New York (2007)

Winikoff, M., Padgham, L., Harland, J., Thangarajah, J.: Declarative and procedural goals in intelligent agent systems. In: Proc. 8th International Conference on Principles of Knowledge Representation and Reasoning (KR), pp. 470–481. Kaufmann, Los Alamitos (2002)

Wooldridge, M., Ågotnes, T., Dunne, P.E., van der Hoek, W.: Logic for automated mechanism design—a progress report. In: Proc. 22nd National Conference on Artificial Intelligence (AAAI), p. 9. AAAI Press, Menlo Park (2007)