

Light-Weight Intelligent Software Agents in Simulation¹

October 1999

Technical Report 99-03

By:

Paolo Busetta, Ralph Rönquist, Andrew Hodgson,
Andrew Lucas
Agent Oriented Software Pty. Ltd. Melbourne, Australia
{paolo,ralph,ash,cal@agent-software.com}
www.agent-software.com

¹ This article first appeared in the proceedings of the 4th International SimTecT Conference '99 pp 169-174. It has been edited for clarity and updated to reflect the current JACK Intelligent Agents distribution.

Abstract

Today Intelligent Agents are being used for modelling human behaviours in simulation. In particular, powerful Belief-Desire-Intention agents have been used successfully in air operations simulations where modelling of human reasoning has been needed to simulate tactical decision making and command and control structures.

However, Intelligent Agent frameworks have so far been large, monolithic software systems. With their origins in research on Distributed Artificial Intelligence, these frameworks have generally been developed as research environments in the research laboratory. Consequently they have been unduly complex to use, as well as expensive and difficult to integrate as components of larger systems.

The JACK Intelligent Agents™ framework presented in this paper brings the concept of intelligent agents into the mainstream of software engineering. JACK is a third generation agent framework, designed to be light-weight, with high performance, and strong typing. It merges the Intelligent Agent concept with the object-oriented Java language, and is designed for building agents for simulation.

We discuss the advantages of having Agents as light-weight components for simulation, and of defining these agents in a mainstream object-oriented language.

Keywords: Java, Intelligent Agents, BDI reasoning

1 Introduction

Intelligent Agents are being used for modelling simple rational behaviours in a wide range of distributed applications, including simulation. There are however, various, if not contradictory, definitions of intelligent agent. By general consensus, it is a type of software that shows some degree of autonomy and social ability, and combines pro-active and reactive behaviours (Wooldridge & Jennings 1995). One of the better known and most successful paradigms for intelligent agents is the BDI (Belief-Desire-Intention) architecture, which has seen a number of academic and industrial applications, including military simulation of both independent and cooperative entities.

Agent Oriented Software Pty. Ltd. (AOS), based in Melbourne, Australia, has built JACK Intelligent Agents™², which is a framework in Java for multi-agent system development. The company's aim is to provide a platform for commercial, industrial and research applications; one of the most important is simulation, in particular, for defence and other domains characterised by a large number of entities showing sophisticated behaviour. JACK supplies a high performance, light-weight implementation of the BDI architecture as an extension of its host language, Java, and can easily support additional agent models and specific application requirements.

In the following, we discuss intelligent agents in the context of simulation. We introduce JACK and present the BDI architecture it currently supports. We give an overview of application development with JACK and conclude with an analysis of the benefits it provides.

2 Intelligent Agents in Simulation

Some of the more advanced applications of computer simulation are in defence and important uses of intelligent agents have been seen in this domain. The considerations reported below arise mostly from experience in military simulation. However, they apply equally to other simulation domains, such as emergency systems (firefighting, natural disaster protection and so on), and any activity involving costly equipment, high levels of training and standard operating procedures (for instance, commercial aviation and air-traffic control).

Historically, simulation in defence has been used for the evaluation of acquisitions or force development options (Lucas & Goss 1999). Modelling and simulation for this purpose is becoming increasingly complex as the number of entities grows and sophistication of scenarios increases. The complexity of this task is further increased by the difficulty in modelling human decision-making with sufficient fidelity using conventional software approaches. Early applications of intelligent agents to represent operational military reasoning have proved highly effective. This success comes from the capability of agents to represent individual reasoning, and from the architectural advantages of that representation to the user due to the ease of setting up and modifying operational reasoning or tactics for various studies.

In addition, certain classes of agents (in particular BDI, discussed later) extend the modelling of reasoning to explicitly model the communications between simulated entities and the coordination of joint activities required for team behaviour. This is particularly helpful when conducting studies that take into consideration complex command and control systems. These are difficult to analyse manually, and advanced modelling and simulation tools are required. Intelligent agents can represent the reasoning and command capabilities appropriate to the humans' assigned roles in the hierarchy, allowing different command and control strategies to be rapidly evaluated under varying circumstances.

The high cost of live exercises has required the development of realistic synthetic training environments. However, so far these synthetic environments have not been able to model the behaviour of the humans involved, other than in a very simple manner. In particular, they have not modelled team behaviour, with the result that trainees quickly learn the range of simulated behaviours: rather than practicing their skills, they learn to predict the training system's response. Intelligent agents allow the simulated humans in training systems to behave in a more credible manner, with a much richer set of behaviours including team responses and dynamic role re-allocation. The result is a more effective training environment

² For the sake of brevity, we refer to JACK Intelligent Agents simply as "JACK".

with realistic tactical behaviour represented, whilst avoiding the expense of having humans involved to provide this.

2.1 A case for light-weight agents

The increasing complexity of equipment and command and control structures, together with the growing dependency on high quality communications (both voice/image and digital) in most human activities, has lead to a corresponding growth in complexity of the software for modelling and simulation. This is further complicated by the move from modelling small numbers of entities involved in short-term engagements to modelling complete theatres of war over several months. The consequent requirements in terms of computational power are growing exponentially, due to increases in communications, in the amount of information to be handled, and in the sophistication of the behaviours of the simulated entities and their interactions (e.g., teams).

Simulation systems that perform satisfactorily for relatively small scenarios do not necessarily scale well; in fact, intelligent agent behaviour is potentially a problematic area, since it has been based on technology more appropriate to AI laboratory research than large-scale deployment. Even when considering the expected gains in computing power and network speed (which historically have been shown to be linear over time), we believe that light-weight implementations of agents are needed in order to match the exponential growth in requirements.

An issue relevant to this discussion is the rapid obsolescence of computing platforms. As a consequence, the choice of a technology for implementing agents should be made so that it is possible to benefit from continuing rapid enhancements to hardware and software systems.

3 JACK Intelligent Agents

In this section, we present JACK by first highlighting the goals set by its designers, then by providing an overview of the engineering characteristics of the framework and finally by describing some features of interest to simulation systems.

3.1 Approach

Major design goals for JACK were: to provide developers with a robust, stable, light-weight product; to satisfy a variety of practical application needs; to ease technology transfer from research to industry; and to facilitate further applied research.

Whilst applications can be built from the ground up adopting an agent oriented methodology and an appropriate framework, most organisations already possess and depend upon large legacy software systems. This is particularly true in the area of training, where large investments into sophisticated HCI systems may have been done over time. Thus, JACK has been designed primarily for use as a component of larger environments. Consequently, an agent must coexist and be visible as simply another object by non-agent software. Conversely, a JACK programmer must be allowed to easily access any other component of a system. Type safeness when accessing data, reliability and support for a proper engineering process are then key requirements in this kind of environment.

For similar reasons, JACK agents are not bound to any specific agent communications language. Instead, they are geared for industrial object-oriented middleware (such as CORBA) and message passing infrastructures (such as DIS). In addition, JACK provides a native light-weight inter-agent communications infrastructure for situations where high performance is paramount.

JACK has been designed for extension by properly trained engineers, familiar with agent concepts and with a sound understanding of concurrent object-oriented programming. The choice of Java as development platform was made after considering: its availability on all modern and foreseeable future platforms; the high quality of its implementation; the rapidly growing body of off-the-shelf software components and interfaces to external systems (such as databases); and its increasing acceptance by both the marketplace and computer practitioners.

3.2 Overview of the Framework

From an engineering perspective, JACK consists of architecture-independent facilities, plus a set of *plug-in* components that address the requirements of specific agent architectures. The plug-ins supplied with JACK include support for the BDI model.

To an application programmer, JACK currently consists of three main extensions to Java. The first is a set of syntactical additions to its host language. These additions, in turn, can be broken down as follows:

- a small number of keywords for the identification of the main components of an agent (such as *agent*, *plan* and *event*);
- a set of statements for the declaration of attributes and other characteristics of the components (for instance, the information contained in beliefs or carried by events). All attributes are strongly typed;
- a set of statements for the definition of static relationships (for instance, which plans can be adopted to react to a certain event);
- a set of statements for the manipulation of an agent's state (for instance, additions of new goals or sub-goals to be achieved, changes of beliefs, or interaction with other agents).

Importantly, the programmer can also include Java statements within the components of an agent.

For the convenience of programmers, in particular those with a background in AI, JACK also supports logical variables and cursors. These are particularly helpful when querying the state of an agent's beliefs. Their semantics are mid-way between logic programming languages (with the addition of type checking Java-style) and embedded SQL.

The second extension to Java is a compiler that converts the syntactic additions described above into pure Java classes and statements that can be loaded with, and be called by, other Java code. The compiler also partially transforms the code of plans in order to obtain the correct semantics of the BDI architecture.

Finally, a set of classes (called the *kernel*) provides the required run-time support to the generated code. This includes:

- the automatic management of concurrency among tasks being pursued simultaneously (*intentions* in the BDI terminology);
- the default behaviour of the agent in reaction to events, failure of actions and tasks, and so on; and
- a native light-weight, high performance communications infrastructure for multi-agent applications.

Importantly, the JACK kernel supports multiple agents within a single process. This is particularly convenient for saving system resources. For instance, agents which perform only short computations or share most of their code or data can be grouped together.

A JACK programmer can extend or change the architecture of an agent by providing new plug-ins. In most cases this simply means overriding the default Java methods provided by the kernel or supplying new classes for run-time support. However, it is possible to add further syntactic extensions to be handled by the JACK compiler.

Similarly, a different communications infrastructure can be supplied by overriding the appropriate run-time methods.

Future versions of JACK will extend the base BDI model with new plug-ins and will add a number of development and monitoring tools.

3.3 Support for Simulation

Some of the features of JACK have been designed to explicitly address, or have been influenced by, simulation requirements.

Agents are equipped with a *clock* which can be used for time-driven simulation. The clocks can be configured to measure actual time as well as simulated time; the latter can be accelerated or retarded while running by means of a graphical interface. This allows us, for instance, to debug a model, or to demonstrate or rehearse scenarios by fast-forwarding up to a critical point, and then time-stepping from then on. A clock can even be replaced by an application-specific timer, provided, for instance, by an external source.

A component called API Builder (APIB) is provided to simplify the task of interfacing JACK agents with legacy systems and external sources of data. The programmer defines the format of application-specific data structures with the APIB data definition language. These declarations are then converted into Java and C++ classes that support platform-independent coding and decoding, sending as messages, and storage and retrieval from files and databases.

Significant events and messages being exchanged among agents can be easily traced, analysed by means of the supplied tool for interaction diagrams or by a user-supplied monitor, and logged for off-line analysis.

Future versions will provide additional tools to help build and monitor agents. They will include a graphical development environment, and ways to trace and view individual and joint tasks (or *intentions* in BDI terminology) being pursued by a team of agents. These tools will allow a non-computer savvy user to model and analyse the behaviour of simulated intelligent entities.

4 Belief-Desire-Intention Agents

The BDI agent model supported by JACK has its roots in both philosophy and cognitive science, and in particular in the work of Bratman (1987) on rational agents. A rational agent has bounded resources, limited understanding and incomplete knowledge of what happens in the environment in which it lives. Such an agent has *beliefs* about the world and *desires* to satisfy, driving it to form *intentions* to act. An intention is a commitment to perform a *plan*. In general, a plan is only partially specified at the time of its formulation since the exact steps to be performed may depend on the state of the environment when they are eventually executed. The activity of a rational agent consists of performing the actions that it intended to execute without any further reasoning until it is forced to revise its own intentions by changes to its beliefs or desires. Beliefs, desires and intentions are called the *mental attitudes* (or mental states) of an agent.

Note that BDI agents depart from purely deductive systems and other traditional AI models because of the concept of intentionality, which significantly reduces the extent of deliberation required. The BDI model has demonstrated its suitability to modelling certain types of behaviour, such as the application of standard operational procedures by trained staff. It has been successfully adopted in fields as diverse as simulation of military tactics, application of business rules in workflows, and diagnostics in telecommunications networks.

Based on previous research and practical application, Rao and Georgeff (1992) have described a computational model for a generic software system implementing a BDI agent. Such a system is an example of an event-driven program. In reaction to an event, for instance a change in the environment or its own beliefs, a BDI agent adopts a *plan* as one of its intentions. Plans are pre-compiled procedures that each have a set of conditions that are used to determine their applicability. The process of adopting a plan as one of the agent's intentions may require a selection from a set of candidate plans.

The agent executes the steps of the plans that it has adopted as intentions until further deliberation is required; this may happen because of new events or the failure or successful conclusion of existing intentions.

A step in a plan must consist of: adding a *goal* (that is, a desire to achieve a certain objective) to the agent itself; changing its beliefs; interacting with other agents; or any other atomic action on the agent's own state or the external world.

This abstract BDI architecture has been implemented in a number of systems. Of these, two are of particular relevance to JACK as they represent its immediate predecessors. The first generation is typified by the Procedural Reasoning System (PRS) (Georgeff & Ingrand (1989),

developed by SRI International in the mid '80s. dMARS (D'Inverno et al. 1998), built in the mid '90s by the Australian Artificial Intelligence Institute in Melbourne, Australia, is a second generation system. dMARS has been used as a development platform for a number of technology demonstrator applications, including simulations of tactical decision-making in air operations and air traffic management.

The BDI architecture provides a powerful means for implementing complex, distributed systems using multiple BDI agents. The SWARM air mission model (Tidhar et al. 1995) was the first application to introduce BDI agents in teams. Current Australian research into teams is directed towards BDI agents for command and control (Tidhar et al. 1998).

5 Application Development With JACK

In an ideal setting, a developer building an application with JACK would start by identifying the distributed functional components of the system. The design of a multi-agent application requires a sound understanding of distributed system development and distributed AI principles that we cannot discuss here. However, it should be noted that in practical situations the decision as to how to distribute functionality may be dictated by a number of external constraints, such as the existence of legacy systems (for instance, HCI equipment) or a specific communications infrastructure (such as DIS).

For the sake of this discussion, let us assume that the functionality to be provided by an agent has been identified and that the BDI model has been chosen as appropriate to the task. At this stage, two main activities have to be performed (not necessarily in the order given below):

- Identification of the elementary classes (that is, abstract data types and the operations allowed on them) that are required to manipulate the resources used by the agent. These could be external (relational databases, an HCI device, a GUI and so on) as well as internal (for instance, specific mathematic data structures to represent spatial information).
- Identification of those elements that constitute the mental states of the agent. This involves finding:
 - which external events drive the agent (including messages from other agents);
 - which goals the agent can set for itself;
 - which beliefs influence the adoption of plans; and
 - the procedures (that is, the plans in BDI terms) required to accomplish tasks, achieve goals and react to events in the various possible contexts.

The same process is performed to identify the joint mental attitudes of a team.

In the current version of JACK, the implementation of an agent is a mix of normal Java code for the elementary classes and extended Java, as described in Section 3.2, for the agent-specific components. The plans of a JACK agent are, in general, sequences of operations on elementary objects, manipulations of the mental states (such as submitting sub-goals or changing beliefs) and interactions with other agents.

In order to improve the accessibility of JACK to the analyst not trained in Java or not interested in low-level programming details, future versions will include a graphical development environment. The users will be able to mix text and GUI based development as appropriate to their competence or the tasks at hand.

5.1 An Example

To give a sample of the code of a JACK agent, we have extracted an example from one of the tutorials that are part of the JACK documentation set. The purpose is not to show agent programming but to illustrate how JACK code looks as a straightforward Java extension. This section can be passed over by those not familiar with Java.

This example has agents that “ping” each other, that is, exchange empty messages. The message being exchanged is represented as an *event* that originates with one agent and is received by another:

```
event PingEvent extends MessageEvent {
    int value;

    #posted as ping(int value)
    {
        this.value = value;
    }
}
```

Note the `event` keyword, in place of `class` in Java. The event is also declared as a `MessageEvent`, which means that it can be sent to another agent; this causes JACK to bring in all the required communications support. The `#posted as` statement declares how the event is generated (in this case, by invoking a Java method `ping()` with one integer parameter).

The following plan handles the notification of the event above and replies to its sender by “bouncing” the event back. This simple example is not sensitive to the event context, i.e., there is no restriction on the state of the beliefs of the agent for its applicability; context checking can be introduced as an additional method of the plan.

```
plan BouncingPlan extends Plan {

    #handles event PingEvent ev;
    #sends event PingEvent pev;

    body()
    {
        @send(ev.from, pev.ping(ev.value + 1));
        /// Reply to the sender of the event
    }
}
```

A trivial “ping agent” is defined below. It has a single plan and handles a single event. When its method `ping (String other)` is called, it notifies the `PingEvent` with value 1 to the agent `other`. If the latter is another `PingAgent`, then `BouncingPlan` above is invoked, a `PingEvent` with value 2 is sent back, and so on to infinity.

```
agent PingAgent extends Agent {

    #handles event PingEvent;
    #uses plan BouncingPlan;

    #posts event PingEvent pev;

    void ping (String other)
    {
        send(other, pev.ping(1));
    }
}
```

The application can instantiate as many `PingAgent` agents as it desires. The name of the agents and their network addresses are determined by the communications mechanism in use; as stated before, JACK provides a high performance messaging system with a simple naming scheme.

5.2 Current Applications of JACK

To date, simulation applications using JACK have been built or are being developed, mainly by DSTO in Australia. Cross and Rönquist (1999) describe a simple air combat simulator, built with the aim of allowing comparisons with SWARMM (Tidhar et al.1995). Other projects currently being undertaken include a close action environment simulator for land-based simulation, and development of team-based agent systems for modelling military command and control architectures. Consideration is also being given to simulating maritime operations, as well as tactical air defence.

A demonstrator of a multi-agent resource scheduling system for land surveillance operations, known as Collection Plan Management System, has recently been delivered to DSTO's Land Operations Division; this demonstrator was built as a functional component within a CORBA environment.

6 Benefits of JACK

The approach taken by Agent Oriented Software with JACK has a number of advantages in comparison with other agent frameworks coming from the artificial intelligence world and with standard object-oriented architectures.

The adoption of Java guarantees a widely available, well-supported execution environment. In addition to the promise of the language (summarised by the well known slogan "compile once, run everywhere" by Sun Microsystems), we expect that an increasing number of software components, tools and trained engineers will be available in the next few years.

To the AI researcher the adoption of an imperative, relatively low-level language such as Java means losing some of the expressive power offered by frameworks based on logic or functional languages. However, this is compensated for, not only by the universal availability mentioned above, but also by the modular approach of JACK. As stated in the previous sections, most components of the framework can be both tuned and tailored. This makes JACK particularly suited to experimentation with new agent architectures for evaluating novel functionality (new mental attitudes, different semantics, additional types of knowledge bases, and so on), or to study performance characteristics in specific contexts.

Moreover, when compared with frameworks based on traditional AI languages, JACK has several distinctive advantages due both to design choices and to a proper utilisation of the intrinsic characteristics of Java. The most important is strong typing, which reduces the chances of programming errors introduced by simple mis-typing. It also provides a very basic version control by making sure that interfaces are compatible at run-time. Next, and particularly relevant to the current discussion, is performance, which makes the execution speed of agent code written in JACK comparable to a direct implementation in C or C++. Moreover, the ability to package multiple agents within a single computer process can lead to significant savings in terms of system resources and communication.

Last, but not least, JACK is of industrial strength and accessible to a large community of engineers trained in object-oriented programming. Future versions will provide graphical development tools, making JACK accessible also to analysts not trained in Java.

As previously stated, agents are being successfully applied to simulation. In summary, the approach of JACK offers the following advantages when compared with traditional software technologies or other agent implementations:

- context sensitivity and sophisticated semantics of mental attitudes of the BDI architecture;
- ease of integration with legacy systems; and
- scalability, that is, the ability to support a large number of functional entities.

7 Conclusions

We have discussed how to handle the computational problems coming from the increase in size and complexity of simulation scenarios by the adoption of a concept actively researched in the AI world, intelligent agents. At the same time, we emphasised the need for scalability and for light-weight implementations of sophisticated software agents.

JACK Intelligent Agents is a multi-agent framework that extends the Java language. The current version supports the BDI agent model, and its modularity enables extensions and different models to be easily supported. We have given an outline of the software development process with JACK and presented some examples. Future versions will include a graphical development environment that will allow the use of JACK without needing to access Java.

JACK's characteristics make it more suitable to large-scale simulations and integration with existing infrastructure than frameworks based on traditional AI technologies, while still offering the advantages of agent programming.

References

- Bratman, M. E. (1987) *Intention, Plans, and Practical Reasoning*. Harvard University Press, Cambridge, MA (USA).
- Cross, M. & Rönquist, R. 'A Java Agent environment for simulation and modelling'. In *Proceedings of the Fourth International SimTect Conference*, 1999
- D'Inverno, M., Kinny, D., Luck, M. & Wooldrige, M. (1998) 'A Formal Specification of dMARS', *INTELLIGENT AGENTS IV: Agent Theories, Architectures, and Languages*. Eds M. Singh, M. Wooldrige, & A. Rao , LNAI 1365, Springer-Verlag, 1998.
- Georgeff, M. P., & Ingrand, F.F. (1989) 'Decision - Making in an embedded reasoning system', *Proceedings of the International Joint Conference on Artificial Intelligence*, Detroit, MI (USA).
- Lucas, A., & Goss, S. 'The potential for intelligent software agents in defence simulation', presented at IDC'99 Symposium, 8-10 February 1999, Adelaide, Australia.
- Rao, A. S., & Georgeff, M.P. (1992) 'An Abstract Architecture for Rational Agents'. *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR'92)*, eds C. Rich, W. Swartout & B. Nebel, Morgan Kaufmann Publishers.
- Tidhar, G., Selvestrel, M. & Heinze, C. (1995) 'Modelling Teams and Team Tactics in Whole Air Mission Modelling'. In *Proceedings of the Eighth International Conference on Industrial Engineering Applications of Artificial Intelligence Expert Systems*, eds G. F. Forsyth and M. Ali, pp 373-381, Melbourne, Australia. Gordon and Breach Publishers.
- Tidhar, G., Rao, A. & Sonenberg, L. (1998) 'On Teamwork and Common Knowledge'. *Proceedings of the 1998 International Conference on Multi-Agent Systems*, Paris.
- Wooldridge, M., & Jennings, N.R. (1995) 'Intelligent Agents: Theory and Practice', *The Knowledge Engineering Review*, vol. 10, no 12, pp 115-152.