Simulating Rational Goal-Directed Behaviour Using a Logic-Based Programming Language for Multi-Agent Systems

Rafael H. Bordini[1]

Extended Abstract

An essential aspect of autonomous agents is that they must display *proactive* behaviour. Designing such software then requires explicit consideration of the *goals* the agent ought to achieve, and similarly its implementation also needs to be based on explicit representations of such goals. This is part of the reason why the BDI (Belief-Desire-Intention) agent architecture [16, 17] has since the early 90's been the best known architecture for developing software agents. As the BDI notions are also used in "folk psychology" (i.e., how people ordinarily refer to other people's behaviour) it also makes it useful for modelling goal-directed human behaviour [12].

As well as work on an agent architecture based on the BDI notions, much work was done to formalise these notions using modal logic [18]. Based on that work and also on practical implementation (as reactive planning systems) of the BDI architecture such as PRS [8], Rao created a simple, abstract agent programming language called AgentSpeak(L) [15]. Building on the basic constructs of logic programming, it provides an elegant language for the essential features of a BDI-based reactive planning system.

Starting from the initial definition of AgentSpeak, we have worked on various extensions of the language, for example to allow agents programmed in that language to communication using a speech-act based agent communication language [21]. We also did work to relate back the programming language constructs and the interpreter data structures to the modalities of BDI logic using the operational semantics of the language [5]. This is important for ongoing work on formal verification (more on this below).

While that work was mainly theoretical, it served as a basis for the development of a very elaborate, highly customisable platform for developing multi-agent systems called ***Jason***, which was done in joint work with Jomi Hübner, and made available *open source* at `http://jason.sf.net`. That work culminated in the recent publication of a textbook to make the ideas of agent programming using ***Jason*** accessible to wider audiences [4].

Various features of the ***Jason*** platform make it useful for modelling and simulation of social phenomena. Since the initial ideas of using agent-based techniques for modelling and simulating human societies [9] in the early 90's, this area has grown at incredible pace, with social scientists all over the world increasingly having interest in using computer simulation as a research method. However, most of the available tools for social simulation allow only very simple agents (with no cognition) to be used. Prominent researchers in that area have advocated the need for cognitive agents in certain advanced types of social simulation [6]. We are, therefore, in ongoing work [3, 1], incorporating features in ***Jason*** which will make it an even more powerful platform for developing software based on multi-agent systems techniques, but also facilitate its use as a tool for social simulation, in particular for modelling human goal-directed behaviour. Two examples are as follows:

Environments: In ***Jason***, environment models have to be programmed in Java. Needless to say, a more declarative, high-level language would be very useful for social simulation, where models of the environment are typically very important. This was the motivation that led to the development of the ELMS language, described in [13]. That work has recently been extended [14] to allow environment descriptions to have objects containing social norms that are to be observed only within the confines of an environment location, possible where an institution or organisation is situated (similarly to 'please refrain from smoking' or 'keep silence' signs). Another recent development [19] is the integration of ***Jason*** with a well-known approach for developing multi-agent environment based on the "artifact" abstraction [20], which could help in the development of very elaborate (distributed) environments.

Organisations: An important part of agent-oriented software engineering is related to agent *organisations*, which has received much research attention in the last few years. We are currently working on allowing specifications of agent organisations (with the related notions of roles, groups, relationships between groups, social norms, etc.) to be used in combination with ***Jason*** for the programming of individual agents. The particular organisational model we use is called $\mathcal{M}$OISE$^+$ [10]; an initial integration with ***Jason*** is discussed in [11], and available from `http://moise.sf.net`. One of the advantages of the approach is that the organisation specification is available for the agents to access and possibly change at run time.

An important use of logic-based techniques in the context of software development in multi-agent systems (in particular, with many of its application being "dependable systems") is for formal verification. In previous work, we devised model checking techniques for multi-agent systems programmed in AgentSpeak [2]. While in that work we were specifically interested in model checking multi-agent systems programmed in AgentSpeak, in a recent ongoing project, joint with Michael Fisher (see acknowledgements below) we are interested in developing techniques that would allow model checking for a variety of agent-oriented programming langauges [7].

[1] University of Durham, UK, email: r.bordini@durham.ac.uk

ACKNOWLEDGEMENTS

The work on model checking agent programming languages mentioned above is joint with Michael Fisher (University of Liverpool) and supported by EPSRC grant numbers EP/D054788 (Durham) and EP/D052548 (Liverpool).

REFERENCES

[1] Rafael H. Bordini, Antônio Carlos da Rocha Costa, Jomi F. Hübner, Álvaro F. Moreira, Fabio Y. Okuyama, and Renata Vieira, 'MAS-SOC: a social simulation platform based on agent-oriented programming', *Journal of Artificial Societies and Social Simulation*, **8**(3), (Jun 2005). JASSS Forum, <http://jasss.soc.surrey.ac.uk/8/3/7.html>.

[2] Rafael H. Bordini, Michael Fisher, Willem Visser, and Michael Wooldridge, 'Verifying multi-agent programs by model checking', *Journal of Autonomous Agents and Multi-Agent Systems*, **12**(2), 239–256, (Mar 2006).

[3] Rafael H. Bordini and Jomi F. Hübner, 'Agent-based simulation using bdi programming in *jason*', in *Agents, Simulation and Applications*, eds., Adelinde M. Uhrmacher and Danny Weyns, Taylor and Francis, (2008). To appear.

[4] Rafael H. Bordini, Jomi Fred Hübner, and Michael Wooldridge, *Programming Multi-Agent Systems in AgentSpeak Using **Jason***, Wiley Series in Agent Technology, John Wiley & Sons, 2007.

[5] Rafael H. Bordini and Álvaro F. Moreira, 'Proving BDI properties of agent-oriented programming languages: The asymmetry thesis principles in AgentSpeak(L)', *Annals of Mathematics and Artificial Intelligence*, **42**(1–3), 197–226, (September 2004). Special Issue on Computational Logic in Multi-Agent Systems.

[6] Cristiano Castelfranchi, 'The theory of social functions: Challenges for computational social science and multi-agent learning', *Cognitive Systems Research*, **2**(1), 5–38, (April 2001).

[7] Louise A. Dennis, Berndt Farwer, Rafael H. Bordini, Michael Fisher, and Michael Wooldridge, 'A common semantic basis for bdi languages', in *Fifth international Workshop on Programming Multi-Agent Systems (ProMAS-2007), held with AAMAS 2007, May 15th, Honolulu, HI*, eds., Mehdi Dastani, Amal El Fallah Seghrouchni, Alessandro Ricci, and Michael Winikoff, (2007). To appear in Springer's LNCS series.

[8] Michael P. Georgeff and A. L. Lansky, 'Reactive reasoning and planning', in *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI'87), 13–17 July, 1987, Seattle, WA*, pp. 677–682, Manlo Park, CA, (1987). AAAI Press / MIT Press.

[9] *Simulating Society: The Computer Simulation of Social Phenomena*, eds., Nigel Gilbert and Jim Doran, UCL Press, London, 1994.

[10] Jomi Fred Hübner, Jaime Simão Sichman, and Olivier Boissier, 'Using the $\mathcal{M}$OISE$^+$ for a cooperative framework of MAS reorganisation', in *Proceedings of the 17th Brazilian Symposium on Artificial Intelligence (SBIA'04)*, eds., Ana L. C. Bazzan and Sofiane Labidi, volume 3171 of *LNAI*, pp. 506–515, Berlin, (2004). Springer.

[11] Jomi Fred Hübner, Jaime Simão Sichman, and Olivier Boissier, 'Developing organised multi-agent systems using the MOISE+ model: Programming issues at the system and agent levels', *International Journal of Agent-Oriented Software Engineering*, **1**(3/4), 370–395, (2007).

[12] Emma Norling, 'Folk psychology for human modelling: Extending the bdi paradigm', in *3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), 19-23 August 2004, New York, NY, USA*, eds., Nicholas R. Jennings, Carles Sierra, Liz Sonenberg, and Milind Tambe, pp. 202–209. IEEE Computer Society, (2004).

[13] Fabio Y. Okuyama, Rafael H. Bordini, and Antônio Carlos da Rocha Costa, 'ELMS: an environment description language for multi-agent simulations', in *Environments for Multiagent Systems, State-of-the-art and Research Challenges. Proceedings of the First International Workshop on Environments for Multiagent Systems (E4MAS), held with AAMAS-04, 19th of July*, eds., Danny Weyns, H. van Dyke Parunak, Fabien Michel, Tom Holvoet, and Jacques Ferber, number 3374 in Lecture Notes in Artificial Intelligence, pp. 91–108, Berlin, (2005). Springer-Verlag.

[14] Fabio Y. Okuyama, Rafael H. Bordini, and Antônio Carlos da Rocha Costa, 'Spatially distributed normative objects', in *Proceedings of the Workshop on Coordination, Organization, Institutions and Norms in Agent Systems (COIN), held with ECAI 2006, 28th August, Riva del Garda, Italy*, eds., Guido Boella, Olivier Boissier, Eric Matson, and Javier Vázquez-Salceda, (2006).

[15] Anand S. Rao, 'AgentSpeak(L): BDI agents speak out in a logical computable language', in *Proceedings of the Seventh Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'96), 22–25 January, Eindhoven, The Netherlands*, eds., Walter Van de Velde and John Perram, number 1038 in Lecture Notes in Artificial Intelligence, pp. 42–55, London, (1996). Springer-Verlag.

[16] Anand S. Rao and Michael P. Georgeff, 'Modeling rational agents within a BDI-architecture', in *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, eds., James Allen, Richard Fikes, and Erik Sandewall, pp. 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, (1991).

[17] Anand S. Rao and Michael P. Georgeff, 'BDI agents: From theory to practice', in *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95), 12–14 June, San Francisco, CA*, eds., Victor Lesser and Les Gasser, pp. 312–319, Menlo Park, CA, (1995). AAAI Press / MIT Press.

[18] Anand S. Rao and Michael P. Georgeff, 'Decision procedures for BDI logics', *Journal of Logic and Computation*, **8**(3), 293–343, (1998).

[19] Alessandro Ricci, Michele Piunti, L. Daghan Acay, Rafael H. Bordini, Jomi F. Hübner, and Mehdi Dastani, 'Integrating heterogeneous agent-programming platforms within artifact-based environments', in *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008), 12–16 May, Estoril, Portugal*, eds., Lin Padgham, David Parkes, Jörg Müller, and Simon Parsons. International Foundation for Autonomous Agents and Multiagent Systems (`www.ifaamas.org`), (2008). To appear.

[20] Alessandro Ricci, Mirko Viroli, and Andrea Omicini, "'Give agents their artifacts": The A&A approach for engineering working environments in MAS', in *6th International Joint Conference "Autonomous Agents & Multi-Agent Systems" (AAMAS 2007)*, eds., Edmund Durfee, Makoto Yokoo, Michael Huhns, and Onn Shehory, pp. 601–603, Honolulu, Hawai'i, USA, (14–18 May 2007). IFAAMAS.

[21] Renata Vieira, Alvaro Moreira, Michael Wooldridge, and Rafael H. Bordini, 'On the formal semantics of speech-act based communication in an agent-oriented programming language', *Journal of Artificial Intelligence Research*, **29**, 221–267, (June 2007).