The MCAPL Framework including the Agent Infrastructure Layer an Agent Java Pathfinder

Louise A. Dennis

Department of Computer Science, University of Liverpool

Summary

The MCAPL (Model-checking Agent Programming Languages) framework is a suite of tools for building interpreters for agent programming languages and verifying the correctness of programs running in these interpreters using the model checking technique. It consists of the Agent Infrastructure Layer (AIL) toolkit for building interpreters for rational agent programming languages (BDI languages) as introduced by (Rao and Georgeff 1992) and the Agent JavaPathFinder (AJPF) model checker (L. A. Dennis et al. 2012). AJPF extends the JavaPathfinder (JPF) model checker (Visser et al. 2003) to prove Linear Temporal Logic properties of BDI programs. The distribution also contains a number of programming languages implemented in the AIL. Chief among these are Gwendolen (Louise A. Dennis 2017), the EASS variant of Gwendolen that can be used to program hybrid autonomous systems and GOAL (Hindriks et al. 2001).

Formal Verification and Model Checking

Formal verification is essentially the process of assessing whether a specification given in formal logic is satisfied on a particular formal description of the system in question. Model checking (Clarke, Grumberg, and Peled 1999) has proved an extremely successful approach to formal verification. Model checking takes an executable model of the system in question, defining all the model's possible executions, and then checks a logical property against this model (and, hence, against all possible executions).

BDI Agents and Programming Languages

The predominant view of rational agency is that encapsulated within the BDI model, Here, 'BDI' stands for *Beliefs*, *Desires*, and *Intentions*. Beliefs represent the agent's (possibly incomplete, possibly incorrect) information about itself, other agents, and its environment, desires represent the agent's long-term goals while intentions represent the goals that the agent is actively pursuing. There are many different agent programming languages and agent platforms based, at least in part, on the BDI approach and these have been used in a wide range of applications, particularly for hybrid control of autonomous systems (where there is a strong requirement for verification).

The MCAPL framework has enabled research reported in a variety of papers (Louise A. Dennis, Fisher, and Webster 2015,Fisher, Dennis, and Webster (2013),Louise A. Dennis et al. (2014),L. A. Dennis et al. (2010),Louise A. Dennis et al. (2016),L. A. Dennis et

al. (2012),L. Dennis et al. (2015),Louise A. Dennis, Fisher, and Winfield (2015)). Examples from many of these publications appear in the software repository for the MCAPL framework and are credited to the authors of those examples.

The MCAPL Framework can be found on Sourceforge (Louise A. Dennis 2018)

References

Clarke, E. M., O. Grumberg, and D. Peled. 1999. *Model Checking.* MIT Press.

Dennis, L. A., M. Fisher, N. Lincoln, A. Lisitsa, and S. M. Veres. 2010. "Declarative Abstractions for Agent Based Hybrid Control Systems." In *Proc. 8th Int. Workshop on Declarative Agent Languages and Technologies (Dalt)*, 6619:96–111. LNCS. Springer.

Dennis, L. A., M. Fisher, M. Webster, and R. H. Bordini. 2012. "Model Checking Agent Programming Languages." *Automated Software Engineering* 19 (1). Springer:5–63.

Dennis, Louise A. 2017. "Gwendolen Semantics: 2017." ULCS-17-001. University of Liverpool, Department of Computer Science.

———. 2018. "Modelchecking Agent Programming Languages." 2018. https://sourceforge.net/p/mcapl/mcapl_code/ci/master/tree/.

Dennis, Louise A., Jonathan M. Aitken, Joe Collenette, Elisa Cucco, Maryam Kamali, Owen McAree, Affan Shaukat, et al. 2016. "Agent-Based Autonomous Systems and Abstraction Engines: Theory Meets Practice." In *Towards Autonomous Robotic Systems: 17th Annual Conference, Taros 2016, Sheffield, Uk, June 26–July 1, 2016, Proceedings*, edited by Lyuba Alboul, Dana Damian, and M. Jonathan Aitken, 75–86. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-40379-3_8.

Dennis, Louise A., Michael Fisher, and Matt Webster. 2015. "Two-Stage Agent Program Verification." *Journal of Logic and Computation.*

Dennis, Louise A., Michael Fisher, and Alan F. T. Winfield. 2015. "Towards Verifiably Ethical Robot Behaviour." In *AAAI Workshop on Ai and Ethics (1st International Conference on Ai and Ethics)*. Austin, TX.

Dennis, Louise A., Michael Fisher, Nicholas K. Lincoln, Alexei Lisitsa, and Sandor M. Veres. 2014. "Practical Verification of Decision-Making in Agent-Based Autonomous Systems." *Automated Software Engineering.* Springer US, 1–55. https://doi.org/10.1007/s10515-014-0168-9.

Dennis, Louise, Michael Fisher, Marija Slavkovik, and Matt Webster. 2015. "Formal Verification of Ethical Choices in Autonomous Systems." *Robotics and Autonomous Systems*, –. https://doi.org/http://dx.doi.org/10.1016/j.robot.2015.11.012.

Fisher, Michael, Louise A. Dennis, and Matthew P. Webster. 2013. "Verifying Autonomous Systems." *Commun. ACM* 56 (9):84–93.

Hindriks, K. V., F. S. de Boer, W. van der Hoek, and J.-J.Ch Meyer. 2001. "Agent Programming with Declarative Goals." In *Intelligent Agents Vii*, 1986:228–43. LNAI. Springer.

Rao, A. S., and M. P. Georgeff. 1992. "An Abstract Architecture for Rational Agents." In *Proc. International Conference on Knowledge Representation and Reasoning (Kr&R)*, 439–49. Morgan Kaufmann.

Visser, Willem, Klaus Havelund, Guillaume P. Brat, Seungjoon Park, and Flavio Lerda. 2003. "Model Checking Programs." *Automated Software Engineering* 10 (2):203–32.