

Lecture Notes in Artificial Intelligence 2990

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

João Leite Andrea Omicini
Leon Sterling Paolo Torroni (Eds.)

Declarative Agent Languages and Technologies

First International Workshop, DALT 2003
Melbourne, Australia, July 15, 2003
Revised Selected and Invited Papers

Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

Volume Editors

João Leite

Universidade Nova de Lisboa
Faculdade de Ciências e Tecnologia, Departamento de Informática
2829-516 Caparica, Portugal
E-mail: jleite@di.fct.unl.pt

Andrea Omicini

Università di Bologna
Dipartimento di Elettronica, Informatica e Sistemistica
Via Venezia 52, 47023 Cesena, Italy
E-mail: andrea.omicini@unibo.it

Leon Sterling

University of Melbourne
Department of Computer Science and Software Engineering
Victoria 3010, Australia
E-mail: leon@cs.mu.oz.au

Paolo Torroni

Università di Bologna
Dipartimento di Elettronica, Informatica e Sistemistica
Viale Risorgimento 2, 40136 Bologna, Italy
E-mail: paolo.torroni@unibo.it

Library of Congress Control Number: 2004106089

CR Subject Classification (1998): I.2.11, C.2.4, D.2.4, D.2, D.3

ISSN 0302-9743

ISBN 3-540-22124-7 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable to prosecution under the German Copyright Law.

Springer-Verlag is a part of Springer Science+Business Media

springeronline.com

© Springer-Verlag Berlin Heidelberg 2004
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Olgun Computergrafik
Printed on acid-free paper SPIN: 11008521 06/3142 5 4 3 2 1 0

Preface

Agent metaphors and technologies are increasingly adopted to harness and govern the complexity of today's systems. As a consequence, the growing complexity of agent systems calls for models and technologies that promote system predictability and enable feature discovery and verification. Formal methods and declarative technologies have recently attracted a growing interest as a means to address such issues.

The aim of the DALT 2003 workshop was two-fold. On the one hand, we wanted to foster a discussion forum to export such techniques into the broader community of agent researchers and practitioners. On the other hand, we wanted to bring in the issues of real-world, complex, and possibly large-scale agent system design in the perspective of formal methods and declarative technologies.

Thanks to the very high quality of our program committee, we managed to put together a rich program, including three technical sessions and two panel sessions: *The Use of Declarative Programming for Agent-Oriented Software Engineering*, moderated by Leon Sterling and Andrea Omicini, and *Declarative and Logic-Based Technology for Agent Reasoning and Interactions*, organized and moderated by Rafael Bordini and Wiebe van der Hoek, with the participation of five invited panelists.

This book contains the revised and extended versions of the papers presented at the workshop, as well as three invited contributions by leading researchers of the field. It is composed of three parts: (i) software engineering and multi-agent system prototyping, (ii) agent reasoning, BDI logics and extensions, and (iii) social aspects of multi-agent systems.

As an introduction to this first part (and in some sense, to the whole book), Omicini and Zambonelli suggest a new view of MAS as complex systems, inspired by some of the recent results in evolutionary biology, and discuss the many different roles that declarative models and technologies can play in the engineering of complex software systems and of MAS in particular.

Castaldi, Costantini, Gentile and Tocchio discuss how a suitably-designed logic-based agent infrastructure (Lira), combined with a logic-based agent-oriented development framework (DALI), makes it possible to supervise and possibly reconfigure the global behavior of a MAS in a dynamic way.

Bergenti, Rimassa, and Viroli provide a framework for the formalization of autonomous agents, that promotes multiple and layered views over social agent features such as ACL, ontology, and social role.

Clark and McCabe present Go!, a multi-paradigm programming language aimed at agent-based applications. By showing how it can be used to build a simple but nontrivial MAS, they face a wide range of issues, from intra-agent to inter-agent ones.

In the last paper of the first part, Son, Pontelli, Ranjan, Milligan, and Gupta present a brief overview of an agent-based framework for rapid application proto-

typing, called ϕ -log, which is rooted in the specific field of evolutionary biology, and is meant to exploit the power of declarative programming to free evolutionary biologist from the burden of directly managing the mess of the many heterogeneous bio-informatic tools and data available today for her/his work.

The second part of this book is devoted to models of agent rationality. Traditionally, declarative technologies have always played a key role in capturing the notion of a rational agency, and in defining it in a formal and intuitive way. Notably, modal logic has proved to be a very powerful formalism to express classic agent mental categories, such as beliefs, commitments, goals, and intentions, and to extend them with, and reason about, other notions such as interaction, cooperativity, expectations, and ignorance.

This part starts with an invited contribution by van der Hoek and Lomuscio, in which the authors explore the unknown, promoting ignorance to a first class citizen when reasoning in Multi-Agent Systems. Arguing that being able to reason about what agents *ignore* is just as important as reasoning about what agents know, they motivate and define a non-standard multi-modal logic, by means of a sound and complete axiomatization, to represent and reason about ignorance in Multi-Agent Systems.

Ancona and Mascardi define Coo-BDI, an extension of the BDI architecture with the notion of cooperativity. The proposed ability of having agents collaborate by exchanging and sharing plans in a flexible way has great potential in the implementation of interpreters of BDI programming language.

Moreira, Vieira, and Bordini build on their previously presented structural operational semantics to AgentSpeak(L) – a BDI, agent-oriented, logic programming language – with an extension to account for inter-agent communication. While doing so, the authors touch upon the long-lasting problem of the semantics of speech acts.

Trân, Harland and Hamilton take the challenge of extending BDI theories to multi-agent systems in an interactive, dynamic environment, and attempt to readdress the computational grounding problem. They provide a formalism for observations – the only connection between mind and worlds – and expectations – the mental states associated with observations.

Flax offers a domain-based approach to the problem of computationally limited deduction and reasoning. The notion of *restricted entailment*, together with the corresponding modal logic interpretation, is presented to model resource-bounded reasoning of agents.

The third and last part of this book focusses on agent interaction. Since the early days of agent research, great attention has been devoted to the study of interactions. This has been done at various levels: by adopting a coordination perspective, by trying to standardize agent communication languages and protocols, and defining a semantics for them, and at a higher level of abstraction, by defining agent teams, societies, organizations, institutions, possibly incorporating organizational notions such as roles and hierarchies, and deontic notions such as norms and obligations.

In this, declarative and logic-based approaches have often been used to define communication language semantics and interaction protocols, both in mentalistic approaches and in social approaches. Also, logic has often been used to define and give an operational semantics to the coordination of reasoning in multi-agent systems.

This third part is started by an invited contribution, in which Colombetti, Fornara, and Verdicchio discuss the use of commitments to give a social semantics to agent interaction, defining the semantics of Agent Communication Languages in terms of changes in the social relationships between agents. The *social commitments*, which represent such relationships, are taken to be primitive concepts, underlying the social dimension of Multi-Agent Systems.

Vasconcelos focusses on communication among the components of a multi-agent system, proposing a logic to describe global protocols. A simple notation is employed, based on first-order logic and set theory to represent an expressive class of electronic institutions. The paper provides a formal semantics for the constructs introduced and presents a distributed implementation of a platform to enact electronic institutions specified in such a formalism.

Alberti, Gavanelli, Lamma, Mello, and Torroni take a resource sharing problem as a case study to present and analyze a social semantics for agent interaction. The formalism introduced is an extension of logic programming with an abductive interpretation, and it allows one to formally define, in a simple and declarative way, concepts such as fulfillment, violation, and social expectation. The authors show how to use these concepts to verify the correct behavior of agents interacting in a society that defines the interaction protocols allowed.

Finally, Küngas and Matskin present a model of cooperative problem solving. Linear logic is used for encoding agents' states, goals, and capabilities. Linear logic theorem proving is applied by each agent to determine whether the particular agent is capable of solving the problem alone. If no individual solution can be constructed, then the agent may start negotiation with other agents in order to find a cooperative solution. Partial deduction in linear logic is used to derive a possible deal, and plans are extracted from the proofs, determining agents' responsibilities in cooperative solutions.

We would like to take this opportunity to thank the authors who answered our call with high quality contributions, the invited panelists, the panel organizers, and all the workshop attendants, for the deep and stimulating discussions, and the authors of the three invited papers. Finally, we would like to thank the members of the Program Committee for ensuring the quality of the workshop program by kindly offering their time and expertise so that each paper could undergo quadruple reviewing.

March 2004

João Leite
Andrea Omicini
Leon Sterling
Paolo Torroni

Organization

DALT 2003 was held in conjunction with AAMAS 2003, the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems, and in cooperation with the ITC Melbourne.

Workshop Organizers

João A. Leite, Universidade Nova de Lisboa, Portugal
Andrea Omicini, Università di Bologna / Cesena, Italy
Leon Sterling, University of Melbourne, Australia
Paolo Torroni, Università di Bologna, Italy

Program Committee

Rafael H. Bordini (The University of Liverpool, UK)
Jeff Bradshaw (The University of West Florida, FL, USA)
Antonio Brogi (Università di Pisa, Italy)
Stefania Costantini (Università degli Studi di L'Aquila, Italy)
Yves Demazeau (Institut IMAG, Grenoble, France)
Jürgen Dix (The University of Manchester, UK)
Toru Ishida (Kyoto University, Japan)
Catholijn Jonker (Vrije Universiteit Amsterdam, NL)
Antonis Kakas (University of Cyprus, Lefkosia, Cyprus)
Daniel Kudenko (University of York, UK)
Alessio Lomuscio (King's College, London, UK)
Viviana Mascardi (Università degli Studi di Genova, Italy)
Paola Mello (Università di Bologna, Italy)
John Jules Ch. Meyer (Universiteit Utrecht, NL)
Charles L. Ortiz (SRI International, Menlo Park, CA, USA)
Sascha Ossowski (Universidad Rey Juan Carlos, Madrid, Spain)
Luís Moniz Pereira (Universidade Nova de Lisboa, Portugal)
Jeremy Pitt (Imperial College, London, UK)
Ken Satoh (National Institute of Informatics, Tokyo, Japan)
Michael Schroeder (City University, London, UK)
Onn Shehory (IBM Research Lab in Haifa, Israel)
Carles Sierra (Spanish Research Council, Barcelona, Spain)
V.S. Subrahmanian (University of Maryland, MD, USA)
Francesca Toni (Imperial College, London, UK)
Wiebe van der Hoek (The University of Liverpool, UK)
Franco Zambonelli (Università di Modena e Reggio Emilia, Italy)

Additional Referees

Álvaro Freitas Moreira
Pavlos Moraitis
Josep Puyol-Gruart
Sarvapali Ramchurn
Juan A. Rodríguez-Aguilar
Ralf Schweimeier

Marek Sergot
Arnon Sturm
Hans van Ditmarsch
Pinar Yolum

Table of Contents

MAS as Complex Systems: A View on the Role of Declarative Approaches	1
<i>Andrea Omicini and Franco Zambonelli</i>	
A Logic-Based Infrastructure for Reconfiguring Applications	17
<i>Marco Castaldi, Stefania Costantini, Stefano Gentile, and Arianna Tocchio</i>	
Operational Semantics for Agents by Iterated Refinement	37
<i>Federico Bergenti, Giovanni Rimassa, and Mirko Viroli</i>	
Go! for Multi-threaded Deliberative Agents	54
<i>Keith L. Clark and Frank G. McCabe</i>	
An Agent-Based Domain Specific Framework for Rapid Prototyping of Applications in Evolutionary Biology	76
<i>Tran Cao Son, Enrico Pontelli, Desh Ranjan, Brook Milligan, and Gopal Gupta</i>	
A Logic for Ignorance	97
<i>Wiebe van der Hoek and Alessio Lomuscio</i>	
Coo-BDI: Extending the BDI Model with Cooperativity	109
<i>Davide Ancona and Viviana Mascardi</i>	
Extending the Operational Semantics of a BDI Agent-Oriented Programming Language for Introducing Speech-Act Based Communication	135
<i>Álvaro F. Moreira, Renata Vieira, and Rafael H. Bordini</i>	
A Combined Logic of Expectation and Observation: A Generalisation of BDI Logics	155
<i>Bình Vũ Trân, James Harland, and Margaret Hamilton</i>	
A Proposal for Reasoning in Agents: Restricted Entailment	173
<i>Lee Flax</i>	
A Social Approach to Communication in Multiagent Systems	191
<i>Marco Colombetti, Nicoletta Fornara, and Mario Verdicchio</i>	
Logic-Based Electronic Institutions	221
<i>Wamberto W. Vasconcelos</i>	

Modeling Interactions Using *Social Integrity Constraints*:
A Resource Sharing Case Study 243
*Marco Alberti, Marco Gavanelli, Evelina Lamma, Paola Mello,
and Paolo Torroni*

Linear Logic, Partial Deduction and Cooperative Problem Solving 263
Peep K ungas and Mihhail Matskin

Author Index 281