

Lecture Notes in Artificial Intelligence 3487

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

João Leite Paolo Torroni (Eds.)

Computational Logic in Multi-Agent Systems

5th International Workshop, CLIMA V
Lisbon, Portugal, September 29-30, 2004
Revised Selected and Invited Papers



Springer

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, Departamento de Informática
Faculdade de Ciências e Tecnologia
Quinta da Torre, 2829-516 Caparica, Portugal
E-mail: jleite@di.fct.unl.pt

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: 2005929660

CR Subject Classification (1998): I.2.11, I.2, C.2.4, F.4

ISSN	0302-9743
ISBN-10	3-540-28060-X Springer Berlin Heidelberg New York
ISBN-13	978-3-540-28060-6 Springer 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. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11533092 06/3142 5 4 3 2 1 0

Preface

The notion of agency has recently increased its influence in the research and development of computational logic based systems, while at the same time significantly gaining from decades of research in computational logic. Computational logic provides a well-defined, general, and rigorous framework for studying syntax, semantics and procedures, for implementations, environments, tools, and standards, facilitating the ever important link between specification and verification of computational systems.

The purpose of the Computational Logic in Multi-agent Systems (CLIMA) international workshop series is to discuss techniques, based on computational logic, for representing, programming, and reasoning about multi-agent systems in a formal way. Former CLIMA editions were conducted in conjunction with other major computational logic and AI events such as CL in July 2000, ICLP in December 2001, FLoC in August 2002, and LPNMR and AI-Math in January 2004.

The fifth edition of CLIMA was held Lisbon, Portugal, in September 29–30, 2004. We, as organizers, and in agreement with the CLIMA Steering Committee, opted for co-location with the 9th European Conference on Logics in Artificial Intelligence (JELIA 2004), wishing to promote the CLIMA research topics in the broader community of logics in AI, a community whose growing interest in multi-agent issues has been demonstrated by the large number of agent-related papers submitted to recent editions of JELIA.

The workshop received 35 submissions – a sensible increase from the previous edition. The submitted papers showed that the logical foundations of multi-agent systems are felt by a large community to be a very important research topic, upon which classical AI and agent-related issues are to be addressed.

In line with the high standards of previous CLIMA editions, the review process was very selective, the final acceptance rate being below 50%. A Program Committee of 24 top-level researchers from 11 countries and 12 additional reviewers selected 16 papers for presentation, authored by 46 researchers worldwide. The workshop program featured an invited lecture by Alessio Lomuscio (University College London) on Specification and Verification of Multiagent Systems, as well as a panel discussion organized by Marina de Vos (University of Bath) on Logic-Based Multi-agent Systems and Industry. Around 50 delegates attended the two-day event.

This book contains a selection, based on a second round of reviewing, of extended CLIMA V papers, and it starts with an invited contribution by Bożena Woźna and Alessio Lomuscio. The papers are divided into four parts: (i) foundations, (ii) architectures, (iii) interaction, and (iv) planning and applications. There follows a brief overview of the book.

Foundations. In the first paper of this book, *A Logic for Knowledge, Correctness, and Real Time*, Woźna and Lomuscio present and exemplify TCTLKD, a logic for knowledge, correctness and real time interpreted on real-time deontic interpreted systems, and extension to continuous time of deontic interpreted systems.

In *Dynamic Logic for Plan Revision in Intelligent Agents*, van Riemsdijk et al. present, with a sound and complete axiomatization, a dynamic logic for a propositional version of the agent programming language 3APL, tailored to handle the revision of plans.

Grossi et al. present in their paper *Contextual Taxonomies* a characterization of the notion of a taxonomy with respect to specific contexts, addressing problems stemming from the domain of normative system specifications for modelling multi-agent systems.

From Logic Programs Updates to Action Description Updates is where Alferes et al. propose a macro language for the language EVOLP and provide translations from some fragments of known action description languages into the newly defined one.

In *Dynamic Logic Programming: Various Semantics Are Equal on Acyclic Programs*, Homola investigates multi-dimensional dynamic logic programming, establishing some classes of programs for which several known semantics coincide.

Architectures. *Declarative Agent Control*, by Kakas et al., extends the architecture of agents based upon fixed, one-size-fits-all cycles of operation by providing a framework for the declarative specification of agent control in terms of *cycle theories*, which define possible alternative behaviors of agents.

In *Metareasoning for Multi-agent Epistemic Logics*, Arkoudas and Bringsjord present an encoding of a sequent calculus for a multi-agent epistemic logic in Athena, an interactive theorem proving system for many-sorted first-order logic, to enable its use as a metalanguage in order to reason about the multi-agent logic as an object language.

In *Graded BDI Models for Agent Architectures*, Casali et al. propose a general model for a graded BDI agent, specifying an architecture able to deal with the environment uncertainty and with graded mental attitudes.

Interaction. Dastani et al., in their article *Inferring Trust*, extend Liau's logic of Belief, Inform and Trust in two directions: with questions, and with a formalization of topics used to infer trust in a proposition from trust in another proposition.

In *Coordination Between Logical Agents*, Sakama and Inoue investigate on the use of answer set programming for belief representation, namely by addressing the problem of finding logic programs that combine the knowledge from different agents, while preserving some properties, useful to achieve agent coordination.

In *A Computational Model for Conversation Policies for Agent Communication*, Bentahar et al. propose a formal specification of a flexible persuasion proto-

col between autonomous agents, using an approach based on social commitments and arguments, defined as a combination of a set of conversation policies.

The last paper of this section is *Verifying Protocol Conformance for Logic-Based Communicating Agents*, by Baldoni et al., which describes a method for automatically verifying a form of “structural” conformance by translating AUML sequence diagrams into regular grammars and, then, interpreting the problem of conformance as a problem of language inclusion.

Planning and Applications. In the preliminary report *An Application of Global Abduction to an Information Agent Which Modifies a Plan Upon Failure*, Satoh uses a form of abductive logic programming called global abduction to implement an information agent that deals with the problem of plan modification upon action failure.

In *Planning Partially for Situated Agents*, Mancarella et al. use an abductive variant of the event calculus to specify planning problems as the base of their proposal for a framework to design situated agents capable of computing partial plans.

Han and Barber, in *Desire-Space Analysis and Action Selection for Multiple Dynamic Goals*, use macro actions to transform the state space for the agent’s decision problem into the desire space of the agent. Reasoning in the latter allows us to approximately weigh the costs and benefits of each of the agent’s goals at an abstract level.

Hirsch et al. conclude this book with the article *Organising Software in Active Environments*, in which they show how logic-based multi-agent systems are appropriate to model active environments. They do so by illustrating how the structuring of the “agent space” can represent both the physical and virtual structures of an application.

We would like to conclude with a glance at the future of this workshop series. The sixth CLIMA edition is being organized by Francesca Toni and Paolo Torroni, and will take place at the City University of London, UK, in June 27–29, 2005, in conjunction with the EU-funded SOCS Project Dissemination Workshop. CLIMA VI will feature a tutorial program and a competition, besides the usual technical content based on the presentation of papers.

We can not miss this opportunity to thank the authors and delegates, who made of CLIMA a very interesting and fruitful event; our generous Program Committee members who did not skimp on time to help us put together a very rich volume after two rounds of reviewing, discussion, and selection; and our sponsoring institutions, Universidade Nova de Lisboa, Fundação para a Ciência e Tecnologia, FBA, and AgentLink III.

April 2005

João Leite
Paolo Torroni

Organization

Workshop Chairs

João Leite, New University of Lisbon, Portugal

Paolo Torroni, University of Bologna, Italy

Program Committee

José Alferes, New University of Lisbon, Portugal

Gerd Brewka, University of Leipzig, Germany

Jürgen Dix, Technical University of Clausthal, Germany

Klaus Fisher, DFKI, Germany

Michael Fisher, The University of Liverpool, UK

James Harland, Royal Melbourne Institute of Technology, Australia

Katsumi Inoue, National Institute of Informatics, Japan

Sverker Janson, Swedish Institute of Computer Science, Sweden

João Leite, New University of Lisbon, Portugal

Yves Lespérance, York University, Canada

John-Jules Ch. Meyer, Utrecht University, The Netherlands

Leora Morgenstern, IBM, USA

Wojciech Penczek, Polish Academy of Sciences, Poland

Jeremy Pitt, Imperial College London, UK

Enrico Pontelli, New Mexico State University, USA

Fariba Sadri, Imperial College London, UK

Ken Satoh, National Institute of Informatics, Japan

Renate Schmidt, The University of Manchester, UK

Tran Cao Son, New Mexico State University, USA

Francesca Toni, University of Pisa, Italy

Wiebe van der Hoek, The University of Liverpool, UK

Paolo Torroni, University of Bologna, Italy

Makoto Yokoo, Kyushu University, Japan

Cees Witteveen, Delft University of Technology, The Netherlands

Additional Reviewers

Federico Banti

Thomas Eiter

Ulle Endriss

Ullrich Hustadt

Magdalena Kacprzak

Olle Olsson

X Organization

Inna Pivkina
Chiaki Sakama

Kostas Stathis
Maciej Szreter

Gregory Wheeler
Yingqiang Zhang

Secretariat

Filipa Mira Reis

Sílvia Marina Costa

Local Organization

António Albuquerque
Duarte Alvim
Eduardo Barros

Jamshid Ashtari
Joana Lopes
Miguel Maurício

Miguel Morais
Sérgio Lopes

Steering Committee

Jürgen Dix, Technical University of Clausthal, Germany

João Leite, New University of Lisbon, Portugal

Fariba Sadri, Imperial College London, UK

Ken Satoh, National Institute of Informatics, Japan

Francesca Toni, University of Pisa, Italy

Paolo Torroni, University of Bologna, Italy

Sponsoring Institutions



Table of Contents

Foundations

A Logic for Knowledge, Correctness, and Real Time <i>Bożena Woźna, Alessio Lomuscio</i>	1
Dynamic Logic for Plan Revision in Intelligent Agents <i>M. Birna van Riemsdijk, Frank S. de Boer, John-Jules Ch. Meyer</i> ...	16
Contextual Taxonomies <i>Davide Grossi, Frank Dignum, John-Jules Ch. Meyer</i>	33
From Logic Programs Updates to Action Description Updates <i>José Júlio Alferes, Federico Banti, Antonio Brogi</i>	52
Dynamic Logic Programming: Various Semantics Are Equal on Acyclic Programs <i>Martin Homola</i>	78

Architectures

Declarative Agent Control <i>Antonis C. Kakas, Paolo Mancarella, Fariba Sadri, Kostas Stathis, Francesca Toni</i>	96
Metareasoning for Multi-agent Epistemic Logics <i>Konstantine Arkoudas, Selmer Bringsjord</i>	111
Graded BDI Models for Agent Architectures <i>Ana Casali, Lluís Godo, Carles Sierra</i>	126

Interaction

Inferring Trust <i>Mehdi Dastani, Andreas Herzig, Joris Hulstijn, Leendert van der Torre</i>	144
Coordination Between Logical Agents <i>Chiaki Sakama, Katsumi Inoue</i>	161

A Computational Model for Conversation Policies for Agent Communication

Jamal Bentahar, Bernard Moulin, John-Jules Ch. Meyer, Brahim Chaib-draa 178

Verifying Protocol Conformance for Logic-Based Communicating Agents

Matteo Baldoni, Cristina Baroglio, Alberto Martelli, Viviana Patti, Claudio Schifanella 196

Planning and Applications

An Application of Global Abduction to an Information Agent which Modifies a Plan upon Failure - Preliminary Report

Ken Satoh 213

Planning Partially for Situated Agents

Paolo Mancarella, Fariba Sadri, Giacomo Terreni, Francesca Toni . . . 230

Desire-Space Analysis and Action Selection for Multiple Dynamic Goals

David C. Han, K. Suzanne Barber 249

Organising Software in Active Environments

Benjamin Hirsch, Michael Fisher, Chiara Ghidini, Paolo Busetta 265

Author Index 281