

Lecture Notes in Artificial Intelligence

1757

Subseries of Lecture Notes in Computer Science

Edited by J. G. Carbonell and J. Siekmann

Lecture Notes in Computer Science

Edited by G.Goos, J. Hartmanis, and J. van Leeuwen

Berlin
Heidelberg
New York
Barcelona
Hong Kong
London
Milan
Paris
Singapore
Tokyo

Nicholas R. Jennings Yves Lespérance (Eds.)

Intelligent Agents VI

Agent Theories Architectures,
and Languages

6th International Workshop, ATAL'99
Orlando, Florida, USA, July 15-17, 1999
Proceedings

Series Editors

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

Volume Editors

Nicholas R. Jennings
University of Southampton
Department of Electronics and Computer Science
Highfield, Southampton SO17 1BJ, UK
E-mail: nrj@ecs.soton.ac.uk

Yves Lespérance
York University
Department of Computer Science
Toronto, ON, Canada M3J 1P3
E-mail: lesperan@cs.yorku.ca

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Intelligent agents V : agent theories, architectures, and languages ;
6th international workshop ; proceedings / ATAL'99, Orlando, Florida,
USA, July 15-17, 1999. Nicholas R. Jennings ; Yves Lespérance (ed.).
- Berlin ; Heidelberg ; New York ; Barcelona ; Hong Kong ; London ;
Milan ; Paris ; Singapore ; Tokyo : Springer, 2000
(Lecture notes in computer science ; Vol. 1757 : Lecture notes in
artificial intelligence)
ISBN 3-540-67200-1

CR Subject Classification (1991): I.2.11, I.2, C.2.4, D.2, F.3

ISBN 3-540-67200-1 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 for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York
a member of BertelsmannSpringer Science+Business Media GmbH
© Springer-Verlag Berlin Heidelberg 2000
Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP-Berlin, Stefan Sossna
Printed on acid-free paper SPIN: 10719619 06/3142 5 4 3 2 1 0

Preface

Intelligent agents are one of the most important developments in computer science in the 1990s. Agents are of interest in many important application areas, ranging from human-computer interaction to industrial process control. The ATAL workshop series aims to bring together researchers interested in the core aspects of agent technology. Specifically, ATAL addresses issues such as theories of agency, software architectures for intelligent agents, methodologies and programming languages for realizing agents, and software tools for developing and evaluating agent systems. One of the strengths of the ATAL workshop series is its emphasis on the synergies between theories, infrastructures, architectures, methodologies, formal methods, and languages.

This year's workshop continued the ATAL trend of attracting a large number of high-quality submissions. In more detail, 75 papers were submitted to the ATAL-99 workshop, from 19 countries. After stringent reviewing, 22 papers were accepted for presentation at the workshop. After the workshop, these papers were revised on the basis of comments received both from the original reviewers and from discussions at the workshop itself. This volume contains these revised papers.

As with previous workshops in the series, we chose to emphasize what we perceive as an important new trend in agent-based computing. In this case, we were motivated by the observation that the technology of intelligent agents and multi-agent systems is beginning to migrate from research labs to software engineering centers. As the rate of this migration increases, it is becoming increasingly apparent that we must develop principled techniques for analyzing, specifying, designing, and verifying agent-based systems. Without such techniques, agent technology will simply not realize its full potential. Consequently, the ATAL-99 program placed particular emphasis on *agent oriented software engineering* and *the evaluation of agent architectures*. Besides several papers in each of these special tracks, the program also featured two associated panels (organized by Mike Wooldridge and Jörg Müller respectively). Another highlight of this year's program was the invited talks by leading exponents of agent research:

THEORIES John Pollock Rational Cognition in OSCAR

ARCHITECTURES Sarit Kraus Agents for Information Broadcasting

It is both our hope and our expectation that this volume will be as useful to the agent research and development community as its five predecessors have proved to be. We believe that ATAL, and the *Intelligent Agents* series of which these proceedings will form a part, play a crucial role in a rapidly developing field, by focusing specifically on the relationships between the theory and practice of agents. Only through understanding these relationships can agent-based computing mature and achieve its widely predicted potential.

November 1999

Nick Jennings
Yves Lespérance

Workshop Organization

Organizing Committee

| | |
|----------------------|--------------------------------------|
| Nicholas R. Jennings | Queen Mary and Westfield College, UK |
| Yves Lespérance | York University, Canada |

Steering Committee

| | |
|----------------------|--------------------------------------|
| Munindar P. Singh | North Carolina State University, USA |
| Michael Wooldridge | Queen Mary and Westfield College, UK |
| Nicholas R. Jennings | Queen Mary and Westfield College, UK |
| Jörg P. Müller | Siemens Research, Germany |

Program Committee

| | | | |
|-------------------------|---------------|--------------------|---------------|
| Ron Arkin | (USA) | Chitta Baral | (USA) |
| Pete Bonasso | (USA) | Stefan Bussmann | (Germany) |
| Cristiano Castelfranchi | (Italy) | Lawrence Cavedon | (Australia) |
| John-Jules Ch. Meyer | (Netherlands) | Phil Cohen | (USA) |
| Rosaria Conte | (Italy) | Frank Dignum | (Netherlands) |
| Alexis Drogoul | (France) | Jacques Ferber | (France) |
| Klaus Fischer | (Germany) | Michael Fisher | (UK) |
| Stan Franklin | (USA) | Fausto Giunchiglia | (Italy) |
| Piotr Gmytrasiewicz | (USA) | Afsaneh Haddadi | (Germany) |
| Henry Hexmoor | (USA) | Wiebe van der Hoek | (Netherlands) |
| Marc Huber | (USA) | Mark d’Inverno | (UK) |
| David Kinny | (Australia) | Sarit Kraus | (Israel) |
| Michael Luck | (UK) | Jörg P. Müller | (Germany) |
| Martha Pollack | (USA) | Anand Rao | (Australia) |
| Onn Shehory | (USA) | Carles Sierra | (Spain) |
| Munindar Singh | (USA) | Liz Sonenberg | (Australia) |
| Katia Sycara | (USA) | Milind Tambe | (USA) |
| Jan Treur | (Netherlands) | Mike Wooldridge | (UK) |
| Eric Yu | (Canada) | | |

Additional Reviewers

| | |
|---------------------------|-----------------------|
| Rina Azoulay (Schwartz) | Thorsten Bohnenberger |
| Rino Falcone | Luke Hunsberger |
| Jiang Bing Li | Juergen Lind |
| Charles Ortiz | Massimo Paolucci |
| Terry Payne | David V. Pynadath |
| Juan A. Rodriguez-Aguilar | Son Cao Tran |

Table of Contents

Section I: Agent Theories

| | |
|---|----|
| Reasoning about Visibility, Perception and Knowledge | 1 |
| Michael Wooldridge and Alessio Lomuscio | |
| A Spectrum of Modes of Knowledge Sharing between Agents | 13 |
| Alessio Lomuscio and Mark Ryan | |
| Observability-Based Nested Belief Computation for Multiagent Systems and its Formalization | 27 |
| Hideki Isozaki and Hirofumi Katsuno | |
| On the Correctness of PRS Agent Programs | 42 |
| Wayne Wobcke | |
| Incorporating Uncertainty in Agent Commitments | 57 |
| Ping Xuan and Victor R. Lesser | |

Section II: Agent and System Architectures

| | |
|--|-----|
| Rational Cognition in OSCAR | 71 |
| John L. Pollock | |
| Agents for Information Broadcasting | 91 |
| Esther David and Sarit Kraus | |
| On the Evaluation of Agent Architectures | 106 |
| Henry Hexmoor, Marcus Huber, Jörg P. Müller, John Pollock and Donald Steiner | |
| Toward a Methodology for AI Architecture Evaluation | 117 |
| Scott A. Wallace and John E. Laird | |
| Reactive-System Approaches to Agent Architectures | 132 |
| Jaeho Lee | |
| A Planning Component for RETSINA Agents | 147 |
| Massimo Paolucci, Onn Shehory, Katia Sycara, Dirk Kalp, and Anandee Pannu | |
| A Scalable Agent Location Mechanism | 162 |
| Onn Shehory | |

Section III: Agent Languages

| | |
|---|-----|
| Reactivity in a Logic-Based Robot Programming Framework | 173 |
| Yves Lespérance, Kenneth Tam, and Michael Jenkin | |
| Extending ConGolog to Allow Partial Ordering | 188 |
| Chitta Baral and Tran Cao Son | |

| | |
|--|-----|
| Operational Semantics of Multi-agent Organisations | 205 |
| Jacques Ferber and Olivier Gutknecht | |
| Open Multi-agent Systems: Agent Communication and Integration | 218 |
| Rogier M. van Eijk, Frank S. de Boer, Wiebe van der Hoek, and John-Jules Ch. Meyer | |
| Toward Team-Oriented Programming | 233 |
| David V. Pynadath, Milind Tambe, Nicolas Chauvat, and Lawrence Cavedon | |
| Section IV: Agent-Oriented Software Engineering | |
| Agent-Oriented Software Engineering | 248 |
| Stefan Bussmann, Paolo Ciancarini, Keith Decker, Michael Huhns, and Michael Wooldridge | |
| Multiagent System Engineering: The Coordination Viewpoint ... | 250 |
| Paolo Ciancarini, Andrea Omicini and Franco Zambonelli | |
| Using Multi-context Systems to Engineer Executable Agents | 260 |
| Jordi Sabater, Carles Sierra, Simon Parsons, and Nicholas R. Jennings | |
| Structuring BDI Agents in Functional Clusters | 277 |
| Paolo Busetta, Nicholas Howden, Ralph Ronnquist, and Andrew Hodgson | |
| Towards a Distributed, Environment-Centered Agent Framework | 290 |
| John R. Graham and Keith S. Decker | |
| Section V: Decision Making in a Social Context | |
| Variable Sociability in Agent-Based Decision Making | 305 |
| Lisa Hogg and Nicholas R. Jennings | |
| Cooperation and Group Utility | 319 |
| Guido Boella, Rossana Damiano, and Leonardo Lesmo | |
| Relating Quantified Motivations for Organizationally Situated Agents | 334 |
| Thomas Wagner and Victor Lesser | |
| The Role and the Impact of Preferences on Multiagent Interaction | 349 |
| Sviatoslav Brainov | |
| Deliberate Normative Agents: Principles and Architecture | 364 |
| Cristiano Castelfranchi, Frank Dignum, Catholijn M. Jonker, and Jan Treur | |
| Subject Index | 379 |
| Author Index | 381 |

Introduction

Like its predecessors [4, 5, 2, 3, 1], this volume of the Intelligent Agent series focuses on the relationships between the theory and the practice of intelligent autonomous agents. To this end, the volume is divided into five sections, reflecting the major current research and development trends in the field. Section I presents work on agent theories. Section II discusses work on architectures for single agents and architectures for entire systems. Section III deals with a variety of agent languages. Section IV presents work on agent-oriented software engineering. Finally, section V deals with agents making decisions while taking the presence of other agents into account.

Section I: Agent Theories

Wooldridge and Lomuscio address the problem of giving a semantics to attributions of knowledge and perception attitudes that is grounded. In their semantics, the truth conditions of these attitudes are defined in terms of actual states of the agent's internal architecture and their relation to environment states. The logic allows one to model the relation between what is objectively true in the environment, what is visible to the agent, what the agent actually perceives, and what it knows.

Lomuscio and Ryan discuss how modal logic can be used to model how knowledge may be shared between agents in a multi-agent system. They examine various axioms about relations between nested knowledge, i.e. knowledge one agent has about what another agent knows. They classify the logical systems that result when various axioms about such knowledge sharing are adopted, specify the associated semantics, and prove completeness results for a large number of cases.

Isozaki and Katsuno also deal with nested belief, looking more specifically at nested belief change. They propose a representation that handles many cases of nested beliefs and present an efficient algorithm for updating these beliefs taking into account observability, memory, and the effects of action. The algorithm's soundness with respect to the logic is proven.

Wobcke examines how agent programs can be specified and proven correct. He looks specifically at programs based on the PRS architecture. He proposes a methodology for writing such programs and develops a formalism based on dynamic logic and context-based reasoning for proving properties about them.

Xuan and Lesser address the problem of coordination between autonomous agents. In particular, they identify the notion of commitment as a key consideration. By their very nature, commitments between agents involve a significant degree of uncertainty and such uncertainty needs to be taken into account when an agent performs its planning and scheduling activities. To this end, a framework for incorporating uncertainty into commitments is presented and a concomitant negotiation framework for handling such commitments is developed.

Section II: Agent and System Architectures

Pollock discusses his work on rationality and its embodiment in the OSCAR agent architecture. In particular, he focuses on the design of an agent that can make decisions and draw conclusions that are rational when judged against human standards of rationality.

David and Kraus describe the design, implementation and evaluation of an information broadcasting system. In particular, they focus on ways in which agents can characterize users' needs and use these characterizations to ensure information is broadcast in the most efficient manner possible.

Hexmoor et al. report on the panel on the evaluation of agent architectures that took place at the workshop. A number of key issues in this area are raised and then each panel member presents his response.

Wallace and Laird present their ongoing work towards developing a methodology for comparing and evaluating agent architectures. They explore the issues involved in architecture evaluation and identify a number of potential evaluation strategies. Eventually they opt for an approach that involves identifying the fundamental properties required by intelligent agents. Their approach is demonstrated by applying it to the Soar and the CLIPS architectures.

Lee also addresses the problem of evaluating agent architectures. In contrast to Wallace and Laird, he focuses on a particular class of domains: namely, reactive systems, i.e. systems in which an agent must maintain an ongoing interaction with a dynamically changing environment. For this class of system, he identifies a range of necessary features that an agent must possess. The availability (or not) of these features is then used to rate a number of common architectures that have appeared in the literature.

Paolucci et al. describe a planning component for agents that operate in a dynamic environment, have only partial knowledge of this environment, and must cooperate with each other. The component is part of the RETSINA multi-agent architecture. The planner, which uses a hierarchical task network representation, interleaves planning with action execution and monitoring the environment for changes that might invalidate a plan. When a plan cannot be completed due to lack of information, planning may be suspended while information gathering actions are executed.

Shehory addresses the problem of agents finding information and services in large scale, open systems. The problems of the traditional approach to this task (i.e. matchmakers, brokers, yellow pages agents and other forms of middleware agent) are highlighted. Then, an approach is advocated in which each individual agent caches a list of acquaintances. This approach is shown to enable agents to locate one another without the need for middleware agents and to ensure that the associated communication complexity is fairly low.

Section III: Agent Languages

Lespérance et al. present an approach to the development of robot controllers that have "high-level reactivity" using the ConGolog agent programming language. Such controllers maintain a model of the environment and can react to

environmental events or exceptional conditions such as the failure of a robot action to achieve its objectives. They do so by suspending or abandoning the current plan and selecting a new plan that is an appropriate response to the event or exception.

Baral and Son extend the ConGolog agent programming language with a construct that supports hierarchical task networks. This allows plans that involve a partial order over actions to be easily specified. It also bridges the gap between the procedural model of ConGolog and that of rule-based agent languages.

Ferber and Gutknecht show how one can develop a formal semantics for an architectural model of open multi-agent systems based on the notions of role and group. In this model, roles represent functions within a group and an agent can play multiple roles. The semantics formalizes the mechanisms of group creation and admission to a group, and constrains communication between processes belonging to different groups.

Van Eijk et al. extend their work on a programming language for multi-agent systems to support the communication of propositional information between agents and the integration of new agents into an open system. A formal semantics based on transition systems is presented.

Pynadath et al. concentrate on the problem of programming autonomous agents to act as teams. They develop an abstract and domain independent framework for specifying an agent's behavior when it acts as part of a team. This teamwork layer can be added as a form of meta-controller to the individual agents so that they become team-enabled. This approach simplifies and speeds up the process of building cooperating agents since there is a significant amount of design experience and code re-use.

Section IV: Agent-Oriented Software Engineering

Bussmann et al. present a brief review of the issues raised in the workshop's panel on agent-oriented software engineering.

Ciancarini et al. present a more detailed response to the issues raised in this panel. In particular, they highlight the advantages of adopting an approach based on coordination models.

Sabater et al. present a logic-based approach for specifying and implementing intelligent agents. They exploit the notion of multi-context systems to develop modular specifications for a range of agent architectures. They go on to show how such specifications can be made operational through the development of an appropriate execution model.

Busetta et al. describe an approach to agent design that involves the composition of reusable modules called capabilities, which encapsulate related beliefs, events, and plans. Intentions are still posted to a global structure and agents retain control over which to pursue according to the circumstances. The approach is implemented within a Java-based PRS like framework called Jack.

Graham and Decker discuss the internal architecture of the DECAF (Distributed Environment Centred Agent Framework) agent framework. DECAF is a toolkit that supports the design and development of intelligent agents. It pro-

vides a range of built-in facilities that agent designers can exploit in order to rapidly prototype their application. Such features include: communication, planning, scheduling, execution monitoring and coordination.

Section V: Decision Making in a Social Context

Hogg and Jennings tackle the problem of an autonomous agent making decisions in a social setting. They present a formal framework that enables the social implications of an agent's decisions to be assessed. In particular, agents can dynamically vary the degree to which they balance individual utility maximization and social welfare maximization. They then empirically evaluate a range of social decision making functions and show which of them are successful in what sorts of domains.

Boella et al. also address the problem of how to go about making local decisions that have positive benefit for a group of cooperating agents. In particular, they develop a decision theoretic approach to planning that deals with coordination as a team.

Wagner and Lesser present a framework in which an agent's organizational context can be explicitly represented and reasoned about. They show how an agent's knowledge structures need to be extended to incorporate information about the organizational structure in which it is embedded and then indicate how an agent's control regime needs to be modified to take this context into account.

Brainov analyses the impact of an agent's preferences on the types of interactions in which it engages. In particular, the assumption that self-interest is the optimal decision policy for an autonomous agent is challenged. Indeed, a number of instances where self-interested behavior leads to inefficient outcomes are presented.

Castelfranchi et al. discuss the roles that norms can play in a society of autonomous agents. They present a set of principles that define how an agent should behave with respect to both norm adherence and norm violation. The embodiment of these principles within an agent architecture is then detailed and their impact upon an agent's reasoning process is outlined.

November 1999

Nick Jennings
Yves Lespérance