

# Lecture Notes in Artificial Intelligence 1365

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

Munindar P. Singh Anand Rao  
Michael J. Wooldridge (Eds.)

# Intelligent Agents IV

## Agent Theories, Architectures, and Languages

4th International Workshop, ATAL'97  
Providence, Rhode Island, USA  
July 24-26, 1997  
Proceedings



Springer

### Series Editors

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

### Volume Editors

Munindar P. Singh

North Carolina State University, Department of Computer Science  
Raleigh, NC 27695-7534, USA  
E-mail: singh@ncsu.edu

Anand Rao

Australian Artificial Intelligence Institute  
171 La Trobe Street, Melbourne, Victoria 3000, Australia  
E-mail: anand@aaii.oz.au

Michael J. Wooldridge

Queen Mary and Westfield College, University of London  
Department of Electronic Engineering  
London E1 4NS, UK  
E-mail: m.j.wooldridge@qmw.ac.uk

Cataloging-in-Publication Data applied for

**Die Deutsche Bibliothek - CIP-Einheitsaufnahme**

**Intelligent agents IV : agent theories, architectures, and languages ;  
4th international workshop ; proceedings / ATAL'97, Providence,  
Rhode Island, USA, July 24 - 26, 1997. Munindar P. Singh ... (ed.). -  
Berlin ; Heidelberg ; New York ; Barcelona ; Budapest ; Hong Kong  
; London ; Milan ; Paris ; Santa Clara ; Singapore ; Tokyo : Springer,  
1998**

(Lecture notes in computer science ; Vol. 1365 : Lecture notes in  
artificial intelligence)  
ISBN 3-540-64162-9

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

ISSN 0302-9743

ISBN 3-540-64162-9 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 1998  
Printed in Germany

Typesetting: Camera ready by author  
SPIN 10631861 06/3142 - 5 4 3 2 1 0 Printed on acid-free paper

# 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 applying 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.

Seventy-six papers were submitted to the ATAL-97 workshop, from seventeen countries. After stringent reviewing, twenty papers were accepted for full presentation and an additional five for short presentation. After the workshop, these papers were revised on the basis of comments received both from original reviewers and from discussions at the workshop itself. This volume contains these revised papers.

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 clear 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-97 program emphasized *methodologies* for agent systems. Besides several papers on methodologies, the program also featured two panels, one specifically on methodologies and one on agent programming languages. Another highlight of the 1997 program was three invited talks by leading exponents of agent research:

THEORIES	Les Gasser	Theories of Agents and Multi-Agents
ARCHITECTURES	Kurt Konolige	Connecting Software and Physical Agents
LANGUAGES	Danny Lange	Java — Just What Mobile Agents Need?

It is both our hope and our expectation that this volume will be as useful to the agent research and development community as its three predecessors have proved to be. We believe that ATAL and the *Intelligent Agents* series of which this volume is 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 1997

Munindar P. Singh (Raleigh, USA)  
Anand S. Rao (Melbourne, Australia)  
Michael J. Wooldridge (London, UK)

## Workshop Organization

### Organizing Committee

Munindar P. Singh	(GENERAL/AMERICAS CHAIR) North Carolina State University, USA
Anand Rao	(ASIA/PACIFIC-RIM CHAIR) Australian AI Institute, Australia
Michael Wooldridge	(EUROPEAN CHAIR) Queen Mary and Westfield College, UK
David Kinny	(METHODOLOGIES TRACK CHAIR) Australian AI Institute, Australia
Nicholas R. Jennings	Queen Mary and Westfield College, UK
Jörg P. Müller	Zuno Ltd, UK

### Program Committee

Ron Arkin	(USA)	Pete Bonasso	(USA)
Hans-Dieter Burkhard	(Germany)	Cristiano Castelfranchi	(Italy)
John-Jules Ch. Meyer	(Netherlands)	Keith Decker	(USA)
Ed Durfee	(USA)	Jacques Ferber	(France)
Klaus Fischer	(Germany)	Michael Fisher	(UK)
Stan Franklin	(USA)	Fausto Giunchiglia	(Italy)
Piotr Gmytrasiewicz	(USA)	Afsaneh Haddadi	(Germany)
Henry Hexmoor	(USA)	David Kinny	(Australia)
Kurt Konolige	(USA)	Sarit Kraus	(Israel)
Yves Lespérance	(Canada)	James Lester	(USA)
Charles Rich	(USA)	Jeff Rosenschein	(Israel)
Wei-Min Shen	(USA)	Carles Sierra	(Spain)
Devika Subramanian	(USA)	Kurt Sundermeyer	(Germany)
Katia Sycara	(USA)	Milind Tambe	(USA)
Mario Tokoro	(Japan)	Jan Treur	(Netherlands)

### Additional Reviewers

Massimo Benerecetti	Frances Brazier	Bruno Caprile	Rosaria Conte
Marco Daniele	Joeri Engelfriet	Rino Falcone	Petra Funk
C. M. Jonker	Ralf Kühnel	Jürgen Lind	Serafini Luciano
Rina Schwartz	Luciano Serafini	Steven Shapiro	David Tremaine
Pascal van Eck	Gero Vierke		

# Contents

<b>Introduction .....</b>	<b>XI</b>
---------------------------	-----------

## **Section I: Methodologies**

<b>Panel: Methodologies for Multi-Agent Systems .....</b>	<b>1</b>
David Kinny, Jan Treur, Les Gasser, Steve Clark, and Jörg Müller	
<b>Designing Embedded Agents to Optimize End-User Objectives ....</b>	<b>3</b>
Marcel Schoppers and Daniel Shapiro	
<b>Agent Architecture as Object Oriented Design .....</b>	<b>15</b>
Joanna Bryson and Brendan McGonigle	
<b>Implementation of a Cooperative Agent Architecture Based on the Language-Action Perspective.....</b>	<b>31</b>
Egon M. Verharen, Frank Dignum, and Sander Bos	
<b>Toward the Specification and Design of Industrial Synthetic Ecosystems .....</b>	<b>45</b>
Van Parunak, John Sauter, and Steve Clark	

## **Section II: Architectures and Infrastructure**

<b>Bidding Mechanisms for Data Allocation in Multi-Agent Environments .....</b>	<b>61</b>
Rina Schwartz and Sarit Kraus	
<b>Distributed Storage of Replicated Beliefs to Facilitate Recovery of Distributed Intelligent Agents.....</b>	<b>77</b>
Arvind K. Bansal, Kotagiri Ramohanarao, and Anand Rao	
<b>A Customizable Coordination Service for Autonomous Agents ...</b>	<b>93</b>
Munindar P. Singh	
<b>A Behavior-Based Approach to Reactivity and Coordination: A Preliminary Report .....</b>	<b>107</b>
Silvia Coradeschi and Lars Karlsson	

## **Section III: Coordination Planning and Monitoring**

<b>On Explicit Plan Languages for Coordinating Multiagent Plan Execution .....</b>	<b>113</b>
Jaeho Lee and Edmund H. Durfee	
<b>Social Comparison for Failure Detection and Recovery .....</b>	<b>127</b>
Gal A. Kaminka and Milind Tambe	

<b>Multi-Agent Coordination through Coalition Formation</b> .....	143
Onn M. Shehory, Katia Sycara, and Somesh Jha	

#### **Section IV: Formal Methods**

<b>A Formal Specification of dMARS</b> .....	155
Mark d’Inverno, David Kinny, Michael Luck, and Michael Wooldridge	
<b>A Framework for Argumentation-Based Negotiation</b> .....	177
Carles Sierra, Nick R. Jennings, Pablo Noriega, and Simon Parsons	
<b>Agent Modelling in METATEM and DESIRE</b> .....	193
Marco Mulder, Jan Treur, and Michael Fisher	
<b>Semantics for an Agent Communication Language</b> .....	209
Yannis Labrou and Tim Finin	

#### **Section V: Theories**

<b>Formal Semantics for an Abstract Agent Programming Language</b> .....	215
Koen V. Hindriks, Frank S. de Boer, Wiebe van der Hoek, and John-Jules Ch. Meyer	
<b>Intentional Agents and Goal Formation</b> .....	231
Frank Dignum and Rosaria Conte	
<b>A Layered Agent Calculus with Concurrent, Continuous Processes</b> .....	245
Christoph G. Jung and Klaus Fischer	
<b>Approximate Reasoning about Combined Knowledge</b> .....	259
Frédéric Koriche	
<b>On the Epistemic Foundations of Agent Theories</b> .....	275
Ho Ngoc Duc	

#### **Section VI: Architectures and Methodologies**

<b>Facilitating Open Communication in Agent Systems: The InfoSleuth Infrastructure</b> .....	281
Marian H. Nodine and Amy Unruh	
<b>Competition for Attention</b> .....	297
Walter Van de Velde, Sabine Geldof, and Ronald Schrooten	
<b>Analysis and Design of Multiagent Systems Using MAS-CommonKADS</b> .....	313
Carlos A. Iglesias, Mercedes Garijo, José C. González, and Juan R. Velasco	

<b>Multi-Agent Systems as Open Societies — A Design Framework</b>	329
Rune E. Gustavsson	
<b>TKQML: A Scripting Tool for Building Agents</b> .....	339
R. Scott Cost, Ian Soboroff, Jeegar Lakhani, Tim Finin, Ethan Miller, and Charles Nicholas	
<b>Subject Index</b> .....	345
<b>Author Index</b> .....	351



## Introduction

Like its three predecessors [2, 3, 1], this volume of the *Intelligent Agents* series focusses on the relationships between the theory and practice of intelligent autonomous agents. To this end, the volume is divided into six sections, reflecting the major current research and development trends in the intelligent agents field. Section I focusses on *methodologies* for agent systems — principled techniques for designing and implementing agent systems. Section II focusses on *architectures and infrastructures*: the papers in this section describe software architectures and techniques for building effective multi-agent systems. Section III focusses on *coordination*, which has long been recognised as an issue of importance in the multi-agent systems community. Section IV focusses on *formal methods* for agent systems, a particular strength of the ATAL workshop series. Section V focusses on the *theoretical foundations* of agent systems, and finally, Section VI focusses on *architectures and methodologies*.

### Section I: Methodologies

*Marcel Schoppers and Dan Shapiro* develop an approach that gives equal importance to user-centric evaluation criteria on the one hand, and agent-centric or developer-centric criteria on the other. This parity is crucial for developing robust methodologies for designing agent-based systems, especially those applied where user-centric criteria, such as safety, cannot be ignored.

Considering agent construction as a special case of object-oriented design (OOD), *Joanna Bryson and Brendan McGonigle* observes that previous approaches, especially in robotic agents, have usually not been accompanied by methodologies for their application. This has limited their applicability to whatever can be achieved through *ad hoc* construction of agents. Using OOD enables the exploitation of traditional software techniques such as polymorphism and inheritance.

The language-action perspective as developed by Winograd and colleagues applies insights from speech act theory to understand and analyze interactions in human organizations. *Egon Verharen* et al apply the same conceptual approach to develop an architecture of agents in information applications. Their architecture yields the increased flexibility needed for modern applications in open information environments.

*Van Parunak* et al define synthetic ecosystems as multiagent systems consisting of a large number of simple agents that can, however, participate in complex interactions. They argue that synthetic ecosystems are beneficial in several industrial applications, including those to do with manufacturing. They then propose an approach for the design of such systems, including important steps toward a methodology.

### Section II: Architectures and Infrastructure

*Rina Schwartz and Sarit Kraus* consider the problem of allocating data among a number of self-interested data servers. They develop a mechanism for optimal allocation of data in this environment. They model the servers as agents with individual preferences and not subject to any central control. Unlike traditional competitive market pricing, this approach works in settings where each product (data item) has exactly one instance.

*Arvind Bansal et al* apply traditional distributed systems techniques to agent systems. They consider the problem of fault tolerance of agents. Fault tolerance is increasingly important as agents are applied in critical domains. The proposed approach saves the state of each agent with other agents in the system. When the agent recovers from a crash, its state is reconstructed from the dumps taken at the other agents. This approach builds on research into logical clocks and potential causality.

*Munindar Singh* considers the problem of coordinating the actions of agents that are heterogeneous — with differing designs — and autonomous. These agents relinquish their autonomy to some extent in order to coordinate with other agents. The agents are modeled as small skeletons showing their actions or events that are visible to other agents, and which are relevant for coordination. This approach takes specifications in a temporal logic and converts them into guards on the individual events that can then be executed in a distributed manner.

*Sylvia Coradeschi and Lars Karlsson* introduce an approach for building reactive but coordinated agents. Their approach defines behaviors as canned scripts or decision-trees, which are specified along with their expected resources and priority. Agents coordinate in real-time, not through communications, but by identifying their team's tactics and knowing the roles they play in those tactics.

### Section III: Coordination Planning and Monitoring

*Jaeho Lee and Ed Durfee* develop structure circuit semantics (SCS) which makes explicit the semantics of plans, otherwise left implicit in the underlying plan interpreter. It adapts GRAFCET, a formalism of the Petri net family, to capture the operational semantics of plans. This approach analyzes plans of different agents to detect and avoid potential deadlocks.

The robust execution of multiagent plans requires the ability on the part of each agent to detect failures and recover from them. *Gal Kaminka and Milind Tambe* develop an approach for failure detection and recovery based on social comparison theory. Social comparison theory, introduced in the 1950s, involves an agent comparing its behavior to that of other agents in a bid to identify discrepancies or to learn from them. Kaminka and Tambe consider classes of agents, such as the teammates of the given agent, from which the learning can be the most effective.

Coalitions promise an intuitive way to combine agents so as to improve the quality of the tasks they perform. However, determining suitable coalitions is nontrivial. Traditional, game-theoretic, approaches are centralized and intractable. *Onn Shehory et al* propose some well-argued simplifications and extensions with which an alternative approach can be defined. Their approach is practicable for systems involving dozens of agents, which cover many cases of practical interest.

### Section IV: Formal Methods

*Mark d'Inverno et al* formally specify the behaviour of a Procedural Reasoning System (PRS) — a class of systems based on the BDI model of practical reasoning. They provide an abstract model of an idealised system and define the key data structures and the

operations that manipulate these structures. The specification is undertaken in Z, a well-known formal specification language.

While a number of papers in the past have examined particular agent languages, very few have compared different languages with respect to developing the same system. *Marco Mulder et al* perform such a comparative study of two of the well known languages for agent modelling — Concurrent METATEM (a multi-agent programming language based on the METATEM paradigm of executable temporal logic), and DESIRE (a compositional specification framework for intelligent systems). The comparison is carried out with respect to the PRS agent architecture: the authors show how the main features of the PRS family of architectures can elegantly be expressed in both Concurrent METATEM and DESIRE.

*Carles Sierra et al* describe a general framework for negotiation in which agents exchange proposals backed by arguments. These arguments summarise the reasons why the proposals should be accepted. The paper uses a business process management example to illustrate the main concepts of argumentation-based negotiation. The primary contribution of the paper is in providing a formal integration of two well-known threads of reasoning — agent-based negotiation and argumentation-based reasoning.

The Knowledge Query and Manipulation Language (KQML) is a widely used language and protocol for communication between agents. Although the language has been around for a few years, there have to date been no serious attempts at giving a semantic description of it. *Yannis Labrou and Tim Finin* rectify this problem by providing a detailed semantics for three key KQML performatives.

## Section V: Theories

*Koen Hindriks et al* propose an abstract agent programming language with a well defined semantics based on transition systems. The language combines both logic programming and imperative programming constructs. It allows users to write practical reasoning rules that provide a mechanism for goal revision. The approach has both theoretical and practical merit and helps to bridge the gap between theory and practice in this area.

*Frank Dignum and Rosaria Conte* address the issue of autonomous goal formation. The authors consider goal formation through behavioural conformity and goal conformity. An alternative mechanism of goal formation through the adoption of norms is also discussed. General rules for goal formation are formally expressed and applied to the social domain.

*Cristoph Jung and Klaus Fischer* present the COOP calculus, a language for concurrent, continuous inference processes as a means of bridging the gap between theory and practice in hybrid architectures of intelligent agent systems. Term rewriting calculus is used as the basis for specifying the semantics. The language can also be viewed as an extension of the hybrid system INTERRAP used for designing intelligent agents.

*Frédéric Koriche* provides a formal framework for modeling tractable reasoning in resource bounded agents that have very large, inconsistent and uncertain sets of knowledge. The paper proposes a model checking approach and a stepwise procedure for improving approximate answers and allowing their convergence to the correct answer.

*Ho Ngoc Duc* proposes an approach to formalizing resource-bounded agents by combining epistemic logic with complexity analysis. Such an approach considers how long an agent will need to compute the answer to a certain query. The paper once again provides a mechanism for bridging idealistic theories with existing practical agents.

## Section VI: Architectures and Methodologies

*Marian Nodine and Amy Unruh* discuss inter-operability issues in multi-agent systems, with particular reference to the InfoSleuth project at MCC. In particular, they discuss the desirability of using common communication languages with a standardised set of speech acts, and a *service ontology* defining the “nouns” and “adjectives” that agents can refer to. In addition to these basic components, Nodine and Unruh describe *conversations*, which are rather like protocols for cooperation.

*Walter Van de Welde et al* describe the *competition for attention* paradigm for agent-based multi-media applications. The competition for attention paradigm is based on the currently popular paradigm of *push technology*, where users are pro-actively provided with information and other services. However, Van de Welde et al point out that unmodified use of push will almost certainly lead to “information overload”, and that a solution is for agents to *compete* for user’s attention. They describe an implementation of the approach in the WWW site for the Brussels-based Ecran film festival.

*Carlos Iglesias et al* continue the theme of methodologies for agent-based systems developed at previous ATAL workshops, by developing a methodology for agent system development based on the well-known Common KADS methodology for knowledge-based system construction. The methodology results in the generation of seven models, describing the various aspects of the system. The methodology is illustrated by means of a simple travel agency case study. Like Iglesias and colleagues, *Rune Gustavsson* also describes an approach to designing multi-agent systems based on the Common KADS methodology.

*R. Scott Cost et al* describe TKQML — an interpreted programming language for building agent systems based on the TCL (Tool Control Language) scripting language. TQML, as its name suggests, supports messaging in the KQML communication language. TKQML allows rapid prototyping of multi-agent systems, and allows the use of all Tcl/Tk facilities for building GUI front ends.

November 1997

Munindar P. Singh (Raleigh, USA),  
Anand S. Rao (Melbourne, Australia), and  
Michael J. Wooldridge (London, UK)

## References

1. J. P. Müller, M. Wooldridge, and N. R. Jennings, editors. *Intelligent Agents III (LNAI Volume 1193)*. Springer-Verlag: Berlin, Germany, 1995.
2. M. Wooldridge and N. R. Jennings, editors. *Intelligent Agents: Theories, Architectures, and Languages (LNAI Volume 890)*. Springer-Verlag: Berlin, Germany, 1995.
3. M. Wooldridge, J. P. Müller, and M. Tambe, editors. *Intelligent Agents II (LNAI Volume 1037)*. Springer-Verlag: Berlin, Germany, 1996.