

Lecture Notes in Artificial Intelligence 1535

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

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Singapore

Tokyo

Sascha Ossowski

Co-ordination in Artificial Agent Societies

Social Structures and Its Implications
for Autonomous Problem-Solving Agents



Springer

Series Editors

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

Author

Sascha Ossowski
Rey Juan Carlos University, School of Engineering (ESCET)
Campus de Móstoles
Calle Tulipan s/n, E-28933 Móstoles (Madrid), Spain
E-mail: S.Ossowski@escet.urjc.es

Cataloging-in-Publication data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Ossowski, Sascha:

Co-ordination in artificial agent societies : social structures and its implications for autonomous problem solving agents / Sascha Ossowski. - Berlin ; Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ; Singapore ; Tokyo : Springer, 1999
(Lecture notes in computer science ; 1535 : Lecture notes in artificial intelligence)
ISBN 3-540-65495-X

CR Subject Classification (1998): I.2.11, C.2.4, H.5.3

ISBN 3-540-65495-X 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 1999
Printed in Germany

Typesetting: Camera-ready by author
SPIN 10692964 06/3142 – 5 4 3 2 1 0 Printed on acid-free paper

Für meine Eltern

Para Sandra

Foreword

Advances in Computer Science often arise from new ideas and concepts, that prove to be advantageous for the design of complex software systems. The conception of multi-agent systems is particularly attractive, as it promises *modularity* based on the conceptual speciality of an agent, as well as *flexibility* in their integration through appropriate interaction models. While early systems drew upon co-operative agents, recent developments have realised the importance of the notion of autonomy in the design of agent-based applications. The emergence of systems of autonomous problem-solving agents paves the way for complex Artificial Intelligence applications that allow for *scalability* and at the same time foster the *reusability* of their components.

In consequence, an intelligent multi-agent application can be seen as a collection of autonomous agents, usually specialised in different tasks, together with a social model of their interactions. This approach implies a dynamic generation of complex relational structures, that agents need to be knowledgeable of in order to successfully achieve their goals. Therefore, a multi-agent system designer needs to think carefully about conceptualisation, representation and enactment of the different types of knowledge that its agents rely on, for individual problem solving as well as for mutual co-ordination. A *knowledge-oriented* approach to agent system design will promote the emergence of methodologies and tools for building multi-agent applications, thereby providing evidence for the added value that agent technology provides over traditional approaches to the design of intelligent software systems.

The proposal that Sascha Ossowski presents in this book applies the above paradigm to artificial agent societies. A comprehensive survey of the different approaches to co-ordination in societies of (human and artificial) agents is given first. Setting out from a critical analysis of the state of the art, he suggests structuring multi-agent applications in line with a mechanism that he calls Structural Co-operation. Agents are provided with expertise about their environment in order to detect and overcome specific types of problem situations, make use of their social knowledge to mutually adjust their activities, and are coerced towards a coherent collective behaviour through normative rules. The proposed model is formalised theoretically within Game Theory, realised practically by means of an agent architecture, and assessed experimentally by building a prototype of a distributed decision support system for road traffic management and comparing it to an alternative model based on a centralised architecture.

The present work is particularly interesting, as it not only promotes a well-founded formal model of co-ordination in artificial agent societies, but also puts it into practice by means of an operational software architecture, organised as a society of intelligent agents, so as to attack an important real-world problem. In spite of focusing on a particular class of domains and settings, the book constitutes a significant contribution, as it proves the feasibility – in design and operation – of building intelligent software applications as artificial agent societies. The societies of autonomous problem-solving agents, that the book advocates for, are an important step towards a modern approach to intelligent system design, where advanced applications coping with complex tasks are configured out of libraries of reusable autonomous agents on the basis of coercive interaction patterns.

In our group, we are happy to have contributed to the research reported in this book. The reader will encounter interesting new ideas that are likely to find their way into future design methodologies for agent-based intelligent systems.

Madrid, September 1998

José Cuena
Dpt. of Artificial Intelligence
Technical University of Madrid

Preface

In recent years the interest in agent technology – and agent-related publications – has risen enormously. There are books covering a wide range of recent research in the field: it is easy to find publications on particular agent theories and their underlying logics, on specific agent architectures and their dynamics, as well as on certain agent languages, their syntax and semantics. Although this huge amount of specialised publications makes available a large body of valuable knowledge and research experience, it may also appear unstructured and rather abstract to the reader. It becomes increasingly hard for an engineer to bridge the gap between particular theories, architectures or languages and their application to real-world problems.

This book is concerned with a broad and, as I believe, essential aspect of agent technology: how to achieve instrumental, co-ordinated behaviour among artificial agents for real-world applications? In the attempt to give an answer to this question, it has been necessary to draw on quite different strands of research. However, having in mind the aforementioned arguments, I have tried to pay special attention to show how these pieces fit together.

The book focuses on the nexus between theory and practice in multi-agent system research: setting out from a class of problems, it covers the whole process from the development of a multi-agent theory to the design of a multi-agent application. It reports how notions rooted in sociology can be adapted and transferred to societies of *artificial* agents; how they can be used to build and formalise a co-ordination mechanism within these societies, and how this mechanism is integrated in an experimental multi-agent system for urban road traffic management.

The history of this book dates back to when I left Germany, endowed with a “Human Capital and Mobility” grant of the European Union, in order to take on a research assistance at the Artificial Intelligence Department of the Technical University of Madrid (UPM). It compiles the essence of four years of research work that I carried out as a member of the Intelligent Systems research group at UPM, and in the course of which I defended my dissertation. During this time I had the chance to meet many people in different parts of the world, who have supported the realisation of my ideas in one or another way. I would like to thank them all. Still, I am particularly grateful...

— to professor José Cuena, director of the Artificial Intelligence Department of UPM and head of the Intelligent Systems research group, for the attention and the time he dedicated to me. He pushed me to do more than I thought I could. Without him, this book would not have reached its present shape.

— to associate professor Ana García-Serrano, my supervisor during my time at UPM. My esteem for her goes beyond the important contribution she provided for the formation of the work presented in this book. Through all this time, Ana has not just been a “boss”, but also a friend.

— to all colleagues who took the trouble to review draft versions of my work: professor Alfons Crespo (UPV, Spain), professor Yves Demazeau (IMAG, France), professor Manuel Hermenegildo (UPM), assistant professor Josefa Hernández (UPM), associate professor Martín Molina (UPM), Dr. habil. David Pearce (DFKI, Germany), M.Sc. José-Luis Sierra (UPM), associate professor Juan Tejada (UCM, Spain) and many more... Their comments and criticisms did a great job in improving clarity and conciseness of this book. Special thanks to Yves Demazeau and David Pearce for encouraging me to submit the manuscript for publication in the LNAI series of Springer Verlag.

— to all members of the Intelligent Systems research group, for their help with many of the tiresome tasks that have come up in the course of the research that led to this book.

I am particularly indebted to my parents, for their moral and intellectual support. A very special “thank you” is owed to Sandra, for her infinite patience with me and my projects. Without her continuous assistance, encouragement and inspiration this book would not have come to a “happy end”.

Madrid, August 1998

Sascha Ossowski

Contents

1	Introduction.....	1
1.1	Motivation	1
1.2	Scope	7
1.3	Scientific Objectives.....	9
1.4	Terminological Remarks	12
1.5	Organisation of the Book	13
2	Co-ordination	15
2.1	An Informal Characterisation	15
2.2	Analytic Models	16
2.2.1	Organisational Interdependencies	17
2.2.2	Organisational Co-ordination Mechanisms.....	19
2.2.3	Formal and Informal Organisation.....	20
2.3	Formal Approaches	21
2.3.1	Centralised Co-ordination	22
2.3.1.1	Quantitative Models.....	22
2.3.1.2	Qualitative Models.....	22
2.3.2	Decentralised Co-ordination	23
2.3.2.1	The Notion of Games	23
2.3.2.2	Partially Conflicting Games	24
2.3.2.3	Pitfalls of the Co-ordination Process.....	26
2.4	Bottom Line	29
3	Distributed Artificial Intelligence.....	31
3.1	Overview of Research Fields	31
3.1.1	A Short History	32
3.1.2	Autonomous Agents.....	33
3.1.2.1	Characterisation	34
3.1.2.2	Techniques	35
3.1.3	Distributed Problem Solving.....	38
3.1.3.1	Characterisation	39
3.1.3.2	Techniques	40
3.1.4	Multi-Agent Systems	43
3.1.4.1	Characterisation	43
3.1.4.2	Techniques	44

3.2	Co-ordination Mechanisms	48
3.2.1	Organisation.....	48
3.2.2	Multi-Agent Planning	49
3.2.3	Negotiation.....	50
3.2.4	Examples.....	52
3.2.4.1	Shoham and Tennenholtz: Social Laws	52
3.2.4.2	Goldman and Rosenschein: Co-operative Rules.....	53
3.2.4.3	Sichman et al.: Dependence Situations	53
3.2.4.4	Rosenschein and Zlotkin: Rules of Encounter	54
3.3	Co-ordination Frameworks.....	55
3.3.1	Von Martial: Co-ordination of Autonomous Agents' Plans.....	56
3.3.2	Decker: TÆMS and GPGP	57
3.3.3	Lander: the TEAM Framework.....	59
3.4	Bottom Line	61
4	Analysis.....	65
4.1	Limitations of Present Co-ordination Models	65
4.2	The Approach of this Book	68
5	Emergent Co-ordination by Structural Co-operation	71
5.1	The Metaphor: Functionality of Social Structure	71
5.1.1	Structural Functionalism.....	71
5.1.2	A Constructive Perspective	74
5.2	The Problem: Reactive Co-ordination	75
5.2.1	A Formal Model.....	76
5.2.2	The Synchronous Blocks Domain.....	78
5.2.3	An Example Setting	81
5.3	The Mechanism: Structural Co-operation	83
5.3.1	Overview.....	84
5.3.2	The Individual Context	86
5.3.2.1	On the Conception of Goal	86
5.3.2.2	Basic Notions	86
5.3.2.3	Self-Interested Action	87
5.3.3	The Social Context.....	88
5.3.3.1	Basic Notions	88
5.3.3.2	Relations between Plans.....	88
5.3.3.3	Dependence Relations between Agents	90
5.3.3.4	Social Exchange.....	94
5.3.4	The Normative Context.....	95
5.3.4.1	Deontic Notions	96
5.3.4.2	The Effect of Normative Prescriptions	98
5.3.4.3	Social Structure.....	100
5.4	Related Work.....	101
5.5	Bottom Line	103

6 Structural Co-operation and Bargaining.....	105
6.1 Structural Co-operation as a Bargaining Scenario	105
6.1.1 Modelling Co-operation.....	106
6.1.2 Modelling Conflict.....	110
6.1.3 The Associated Bargaining Scenario	112
6.2 Solutions to the Associated Bargaining Scenario.....	113
6.2.1 The Nash Solution.....	113
6.2.2 Alternative Solutions.....	117
6.2.3 Analysis.....	118
6.3 The Role of the Disagreement Point	119
6.4 Related Work.....	121
6.5 Bottom Line	124
7 An Instrumentation of Structural Co-operation.....	127
7.1 The ProsA ₂ Agent Architecture	127
7.1.1 Overview.....	127
7.1.2 Information Model	129
7.1.2.1 Self and Acquaintance Models	130
7.1.2.2 Problem-Solving and Deontic Models	130
7.1.2.3 Individual and Social Models	131
7.1.3 Layer Models	132
7.1.3.1 Individual Layer.....	132
7.1.3.2 Social Layer	133
7.1.3.3 Normative Layer	133
7.2 Social Interaction Strategy	134
7.2.1 Motivation.....	134
7.2.2 Multi-Stage Co-ordination Process	135
7.2.3 Asynchronous Search for Pareto-Optimality	137
7.2.3.1 Overview	138
7.2.3.2 Local Choice	139
7.2.3.3 Knowledge and Information Requirements	140
7.2.3.4 Message Types.....	141
7.2.3.5 Local Behaviour Strategy.....	142
7.2.3.6 Solution Detection	144
7.2.4 Applying the Nash Solution	146
7.2.5 Analysis.....	148
7.3 Related Work.....	150
7.4 Bottom Line	151
8 Road Traffic Management	153
8.1 Urban Traffic Management	153
8.1.1 Traffic Control Infrastructure.....	153
8.1.2 Decision Support for Real-Time Traffic Management	155

8.2	The TRYS Approach.....	156
8.2.1	Motivation.....	157
8.2.2	The Centralised TRYS Architecture	158
8.3	Multi-Agent TRYS.....	159
8.3.1	Motivation.....	160
8.3.2	The Multi-Agent TRYS Architecture	161
9	A Case Study	163
9.1	The Scenario.....	163
9.2	Local Problem Solving.....	165
9.2.1	Conceptual Model	165
9.2.2	Knowledge Model.....	168
9.2.2.1	Physical Network Structure	168
9.2.2.2	Problem Scenarios	169
9.2.2.3	Traffic Distribution Scenarios.....	169
9.2.2.4	Historic Traffic Demand	170
9.2.3	Reasoning.....	170
9.3	Co-ordination	172
9.3.1	Conceptual Model.....	172
9.3.2	Knowledge Model.....	175
9.3.2.1	Plan Interrelations	175
9.3.2.2	Agent Dependence	176
9.3.2.3	Norms.....	177
9.3.3	Reasoning.....	178
9.4	Implementation Issues.....	180
9.4.1	Knowledge Acquisition and Reuse	180
9.4.2	Process Arrangement	181
9.4.3	Object Hierarchy	183
9.5	Evaluation	184
9.5.1	Co-ordination Policy	185
9.5.2	Knowledge Requirements	187
9.5.3	Architecture.....	188
10	Conclusions.....	191
10.1	Contributions	191
10.2	Limitations	193
10.3	Future Work	194
Appendix	197
A.	Proofs.....	197
B.	Examples of Knowledge Bases	201
References	209

List of Figures

Figure 1.1. Agent types	7
Figure 2.1. Interdependencies between tasks and resources.....	18
Figure 2.2. An abstract view of co-ordination.....	19
Figure 2.3. Types of informal structures	21
Figure 3.1. Example of a multi-agent telecom market	47
Figure 3.2. Criteria for the choice of partners	54
Figure 3.3. Plan relations and co-ordination actions	56
Figure 3.4. Example of a TÆMS model.....	58
Figure 3.5. An integrative view of DAI.....	62
Figure 5.1. Emergent co-ordination by structural co-operation.....	75
Figure 5.2. An example scenario.....	82
Figure 5.3. The mechanism of structural co-operation.....	84
Figure 5.4. Degrees of social dependence	93
Figure 6.1. Graphical representation of the example scenario	109
Figure 6.2. Mixed multi-plans and the Nash solution	117
Figure 6.3. Alternative bargaining solutions	118
Figure 6.4. Disagreement point monotonicity	120
Figure 7.1. ProsA ₂ agent architecture	128
Figure 7.2. ProsA ₂ information model.....	130
Figure 7.3. Overview of the social interaction process.....	136
Figure 7.4. Some model revisions in Stage 1	142
Figure 7.5. Asynchronous search for Pareto-optimality	143
Figure 7.6. The solution detection algorithm	146
Figure 7.7. The stages two and three	147
Figure 8.1. Traffic control infrastructure.....	155
Figure 8.2. Centralised TRYSA architecture	159
Figure 8.3. Multi-agent TRYSA architecture.....	161
Figure 9.1. Autonomous traffic agents for Barcelona	164
Figure 9.2. Graphical representation of a problem area	167
Figure 9.3. An example scenario.....	171
Figure 9.4. Relations between local signal plans.....	174
Figure 9.5. TRYSA ₂ layer knowledge	178
Figure 9.6. TRYSA ₂ process arrangement.....	181
Figure 9.7. An example TRYSA ₂ environment	182
Figure 9.8. TRYSA ₂ object diagram.....	183
Figure 9.9. A simulated traffic problem	185

List of Tables

Table 1.1. Distribution and design	6
Table 1.2. Scope of this book	8
Table 2.1. Coherence problem.....	27
Table 2.2. The Prisoners' Dilemma.....	28
Table 5.1. Some example plans and their properties	83
Table 6.1. Scenario matrix	108
Table 9.1. Comparison concerning co-ordination policy	186
Table 9.2. Comparison concerning knowledge requirements.....	187
Table 9.3. Comparison concerning architecture	188