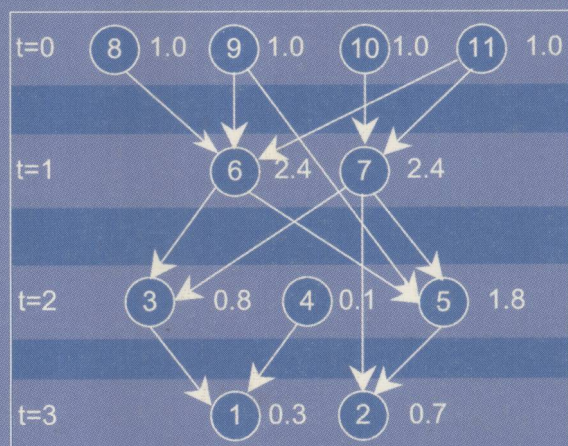


Gabriela Lindemann
Daniel Moldt
Mario Paolucci (Eds.)

Regulated Agent-Based Social Systems

First International Workshop, RASTA 2002
Bologna, Italy, July 2002
Revised Selected and Invited Papers



Springer

TP 3-53
R 344
2002

Gabriela Lindemann Daniel Moldt
Mario Paolucci (Eds.)

Regulated Agent-Based Social Systems

First International Workshop, RASTA 2002
Bologna, Italy, July 16, 2002
Revised Selected and Invited Papers



E200401613



Springer

Series Editors

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

Volume Editors

Gabriela Lindemann
Humboldt-Universität zu Berlin
Institut für Informatik, Lehrstuhl für Künstliche Intelligenz
Rudower Chaussee 25, 10099 Berlin, Germany
E-mail: lindeman@informatik.hu-berlin.de

Daniel Moldt
Universität Hamburg, Fachbereich Informatik
Vogt-Kölln-Straße 30, 22527 Hamburg, Germany
E-mail: moldt@informatik.uni-hamburg.de

Mario Paolucci
Consiglio Nazionale delle Ricerche, CNR
Viale Marx, 15, 00137 Rome, Italy
E-mail: paolucci@ip.rm.cnr.it

Cataloging-in-Publication Data applied for

A catalog record for this book is available from the Library of Congress.

Bibliographic information published by Die Deutsche Bibliothek
Die Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliografie;
detailed bibliographic data is available in the Internet at <<http://dnb.ddb.de>>.

CR Subject Classification (1998): I.2.11, I.2, J.4, C.2.4, I.6

ISSN 0302-9743

ISBN 3-540-20923-9 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 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 PTP-Berlin, Protago-TeX-Production GmbH
Printed on acid-free paper SPIN: 10985007 06/3142 5 4 3 2 1 0

Lecture Notes in Artificial Intelligence 2934

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

Springer

Berlin

Heidelberg

New York

Hong Kong

London

Milan

Paris

Tokyo

Preface

This volume presents selected, extended and reviewed versions of the papers presented at the 1st International Workshop on Regulated Agent Systems: Theory and Applications (RASTA 2002), a workshop co-located with the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002), which was held in Bologna, Italy, in July, 2002. In addition, several new papers on the workshop theme appear here as the result of a further call for participation.

Agent-technology is the latest paradigm of software engineering methodology. The development of autonomous, mobile, and intelligent agents brings new challenges to the field. Agent technologies and multiagent systems are among the most vibrant and active research areas of computer science. At the same time commercial applications of agents are gaining attention. The construction of artificial (agent) societies leads to questions that already have been asked for human societies. Computer scientists have adopted terms like emerging behavior, self-organization, and evolutionary theory in an intuitive manner. Multiagent system researchers have started to develop agents with *social* abilities and complex *social* systems.

However, most of these systems lack the foundation of the *social sciences*. The intention of the RASTA workshop, and of this volume, is to bring together researchers from computer science as well as the social sciences who see their common interest in social theories for the construction and regulation of multi-agent systems.

A total of 17 papers appear in this volume, out of 31 papers submitted. They include nine papers presented in the workshop (whose preproceedings were published as *Communications Vol. 318 Mitteilung 318* of Hamburg University, Faculty of Informatics), as well as six new papers. In addition, an invited paper from Bruce Edmonds reflects some aspects of the lively discussions held during the workshop. The selection presented is divided into two major topics.

Topic A – *Social Theory for Agent Technology (Socionics)*

The wide range of social theories offers many different solutions to problems found in complex (computer) systems. Which theories, and how and when to apply them is a major challenge. In developing agents and multiagent systems computer scientists have used sociological terms like negotiation, interaction, contracts, agreement, organization, cohesion, social order, and collaboration. Meanwhile an interdisciplinary area called socionics, the bridge between sociology and computer science, is beginning to establish itself. The realization that the behavior of societies cannot fully be explained by macrotheories only, and the progress made in agent technology have opened the way to new models of societies in which both macrotheories and microtheories are incorporated. The development

of the socionics research area and the increased interest in the dynamics of the behavior of agents in hybrid organizations requires the investigation of new modelling concepts like roles, groups, social intelligence, emotions, beliefs, desires, and intentions.

Topic B – *Norms and Institutions in MAS*

Multiagent systems are increasingly being considered a viable technological basis for implementing complex, open systems such as electronic marketplaces, virtual enterprises, political coalition support systems, etc. The design of open systems in such domains poses a number of difficult challenges, including the need to cope with unreliable communication and network infrastructures, the need to address incompatible assumptions and limited trust among independently developed agents, and the necessity to detect and respond to systemic failures.

Human organizations and societies have successfully coped with similar problems of coordination, cooperation, etc., in short, with the challenge of social order, mainly by developing norms and conventions, that is, specifications of behavior that all society members are expected to conform to, and that undergo efficient forms of decentralized control. In most societies, norms are backed by a variety of social institutions that enforce law and order (e.g., courts, police), monitor for and respond to emergencies (e.g. ambulance service), prevent and recover from unanticipated disasters (e.g., coast guard, firefighters), etc. In that way, civilized societies allow citizens to utilize relatively simple and efficient rules of behavior, offloading the prevention and recovery of many problem types to social institutions that can handle them efficiently and effectively by virtue of their economies of scale and widely accepted legitimacy. Successful civil societies have thus achieved a division of labor between individuals and institutions that decreases the “barriers to survival” for each citizen, while helping to increase the welfare of the society as a whole.

Several researchers have recognized that the design of open multiagent systems can benefit from abstractions analogous to those employed by our robust and relatively successful societies and organizations. There is a growing body of work that touches upon the concepts of norms and institutions in the context of multiagent systems.

July 2003

Daniel Moldt
Gabriela Lindemann
Mario Paolucci

Organization

The International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA 2002) was organized by: *the Institute of Cognitive Sciences and Technologies* - CNR, Italy; *MIT Sloan School of Management*, USA; *AI Lab of the Department of Computer Sciences*, Humboldt University, Berlin; and *the Theoretical Foundations of Computer Science Group*, University of Hamburg.

Workshop Chairs

Daniel Moldt
Gabriela Lindemann
Mario Paolucci
Bin Yu

Organizing Committee

Rosaria Conte
Chris Dellarocas
Henry A. Kautz
Gabriela Lindemann
Daniel Moldt
Mario Paolucci
Munindar P. Singh
Bin Yu

Program Committee

Andreas Abecker
Karl Aberer
Mark S. Ackerman
Sven Brückner
Kathleen Carley
Jose Carmo
Enhong Chen
Helder Coelho
Rosaria Conte
Noshir Contractor
Raymond D'Amore
Kerstin Dautenhahn

Fiorella De Rosi
Chris Dellarocas
Frank Dignum
Peter Dittrich
Rino Falcone
David Hales
Andrea Hollingshead
Michael Huhns
Andrew Jones
Catholijn Jonker
Henry A. Kautz
Stefan Kirn

VIII Organization

Victor Lesser
Ioan Alfred Letia
Henry Lieberman
Gabriela Lindemann
Jiming Liu
Steve Marsh
Mark Maybury
Ivica Mitrovic
Daniel Moldt
Bonnie Nardi
Hiroaki Ogata
Sascha Ossowski
Pietro Panzarasa
Mario Paolucci
Mirko Petric
Paolo Petta
Michael Prietula

Juan Antonio Rodriguez-Aguilar
Giovanni Sartor
Bernd Schmidt
Ingo Schulz-Schaeffer
Bart Selman
Carles Sierra
Munindar P. Singh
Sorin Solomon
Katia Sycara
Ingo Timm
Inga Tomic-Koludrovic
Adeline Uhrmacher
Thomas Uthmann
Leon Van der Torre
Harko Verhagen
Pinar Yolum
Bin Yu

Referees (*not included in the Program Committee*)

Luis Antunes
Joscha Bach
Francois Bousquet
Jan Broersen
Marc Esteve
Eduardo Fermé
Guido Fioretti
David Hales
Xiaolong Jin

Michael Köhler
Maria Miceli
Dagmar Monett
Tim Norman
Alexander Osherenko
Giovanni Pezzulo
Heiko Rölke
Martijn Schut
Luca Tummolini

Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 2953: K. Konrad, Model Generation for Natural Language Interpretation and Analysis. XIII, 166 pages. 2004.
- Vol. 2934: G. Lindemann, D. Moldt, M. Paolucci (Eds.), Regulated Agent-Based Social Systems. Proceedings, 2002. X, 301 pages. 2004.
- Vol. 2930: F. Winkler, Automated Deduction in Geometry. VII, 231 pages. 2004.
- Vol. 2923: V. Lifschitz, I. Niemelä (Eds.), Logic Programming and Nonmonotonic Reasoning. IX, 365 pages. 2004.
- Vol. 2913: T.M. Pinkston, V.K. Prasanna (Eds.), High Performance Computing - HiPC 2003. Proceedings, 2003. XX, 512 pages. 2003.
- Vol. 2903: T.D. Gedeon, L.C.C. Fung (Eds.), AI 2003: Advances in Artificial Intelligence. Proceedings, 2003. XVI, 1075 pages. 2003.
- Vol. 2902: F.M. Pires, S.P. Abreu (Eds.), Progress in Artificial Intelligence. Proceedings, 2003. XV, 504 pages. 2003.
- Vol. 2892: F. Dau, The Logic System of Concept Graphs with Negation. XI, 213 pages. 2003.
- Vol. 2891: J. Lee, M. Barley (Eds.), Intelligent Agents and Multi-Agent Systems. Proceedings, 2003. X, 215 pages. 2003.
- Vol. 2882: D. Veit, Matchmaking in Electronic Markets. XV, 180 pages. 2003.
- Vol. 2871: N. Zhong, Z.W. Raś, S. Tsumoto, E. Suzuki (Eds.), Foundations of Intelligent Systems. Proceedings, 2003. XV, 697 pages. 2003.
- Vol. 2854: J. Hoffmann, Utilizing Problem Structure in Planning. XIII, 251 pages. 2003.
- Vol. 2843: G. Grieser, Y. Tanaka, A. Yamamoto (Eds.), Discovery Science. Proceedings, 2003. XII, 504 pages. 2003.
- Vol. 2842: R. Gavaldá, K.P. Jantke, E. Takimoto (Eds.), Algorithmic Learning Theory. Proceedings, 2003. XI, 313 pages. 2003.
- Vol. 2838: N. Lavrač, D. Gamberger, L. Todorovski, H. Blockeel (Eds.), Knowledge Discovery in Databases: PKDD 2003. Proceedings, 2003. XVI, 508 pages. 2003.
- Vol. 2837: N. Lavrač, D. Gamberger, L. Todorovski, H. Blockeel (Eds.), Machine Learning: ECML 2003. Proceedings, 2003. XVI, 504 pages. 2003.
- Vol. 2835: T. Horváth, A. Yamamoto (Eds.), Inductive Logic Programming. Proceedings, 2003. X, 401 pages. 2003.
- Vol. 2821: A. Günter, R. Kruse, B. Neumann (Eds.), KI 2003: Advances in Artificial Intelligence. Proceedings, 2003. XII, 662 pages. 2003.
- Vol. 2807: V. Matoušek, P. Mautner (Eds.), Text, Speech and Dialogue. Proceedings, 2003. XIII, 426 pages. 2003.
- Vol. 2801: W. Banzhaf, J. Ziegler, T. Christaller, P. Ditttrich, J.T. Kim (Eds.), Advances in Artificial Life. Proceedings, 2003. XVI, 905 pages. 2003.
- Vol. 2797: O.R. Zaïane, S.J. Simoff, C. Djeraba (Eds.), Mining Multimedia and Complex Data. XII, 281 pages. 2003.
- Vol. 2792: T. Rist, R. Aylett, D. Ballin, J. Rickel (Eds.), Intelligent Virtual Agents. Proceedings, 2003. XV, 364 pages. 2003.
- Vol. 2782: M. Klusch, A. Omicini, S. Ossowski, H. Laamanen (Eds.), Cooperative Information Agents VII. Proceedings, 2003. XI, 345 pages. 2003.
- Vol. 2780: M. Dojat, E. Keravnou, P. Barahona (Eds.), Artificial Intelligence in Medicine. Proceedings, 2003. XIII, 388 pages. 2003.
- Vol. 2777: B. Schölkopf, M.K. Warmuth (Eds.), Learning Theory and Kernel Machines. Proceedings, 2003. XIV, 746 pages. 2003.
- Vol. 2752: G.A. Kaminka, P.U. Lima, R. Rojas (Eds.), RoboCup 2002: Robot Soccer World Cup VI. XVI, 498 pages. 2003.
- Vol. 2741: F. Baader (Ed.), Automated Deduction – CADE-19. Proceedings, 2003. XII, 503 pages. 2003.
- Vol. 2705: S. Renals, G. Grefenstette (Eds.), Text- and Speech-Triggered Information Access. VII, 197 pages. 2003.
- Vol. 2703: O.R. Zaïane, J. Srivastava, M. Spiliopoulou, B. Masand (Eds.), WEBKDD 2002 - Mining Web Data for Discovering Usage Patterns and Profiles. IX, 181 pages. 2003.
- Vol. 2700: M.T. Paziienza (Ed.), Extraction in the Web Era. XIII, 163 pages. 2003.
- Vol. 2699: M.G. Hinchey, J.L. Rash, W.F. Truszkowski, C.A. Rouff, D.F. Gordon-Spears (Eds.), Formal Approaches to Agent-Based Systems. IX, 297 pages. 2002.
- Vol. 2691: V. Mafík, J.P. Müller, M. Pechoucek (Eds.), Multi-Agent Systems and Applications III. Proceedings, 2003. XIV, 660 pages. 2003.
- Vol. 2684: M.V. Butz, O. Sigaud, P. Gérard (Eds.), Anticipatory Behavior in Adaptive Learning Systems. X, 303 pages. 2003.
- Vol. 2671: Y. Xiang, B. Chaib-draa (Eds.), Advances in Artificial Intelligence. Proceedings, 2003. XIV, 642 pages. 2003.
- Vol. 2663: E. Menasalvas, J. Segovia, P.S. Szczepaniak (Eds.), Advances in Web Intelligence. Proceedings, 2003. XII, 350 pages. 2003.
- Vol. 2661: P.L. Lanzi, W. Stolzmann, S.W. Wilson (Eds.), Learning Classifier Systems. VII, 231 pages. 2003.

- Vol. 2654: U. Schmid, Inductive Synthesis of Functional Programs. XXII, 398 pages. 2003.
- Vol. 2650: M.-P. Huget (Ed.), Communications in Multi-agent Systems. VIII, 323 pages. 2003.
- Vol. 2645: M.A. Wimmer (Ed.), Knowledge Management in Electronic Government. Proceedings, 2003. XI, 320 pages. 2003.
- Vol. 2639: G. Wang, Q. Liu, Y. Yao, A. Skowron (Eds.), Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing. Proceedings, 2003. XVII, 741 pages. 2003.
- Vol. 2637: K.-Y. Whang, J. Jeon, K. Shim, J. Srivastava, Advances in Knowledge Discovery and Data Mining. Proceedings, 2003. XVIII, 610 pages. 2003.
- Vol. 2636: E. Alonso, D. Kudenko, D. Kazakov (Eds.), Adaptive Agents and Multi-Agent Systems. XIV, 323 pages. 2003.
- Vol. 2627: B. O'Sullivan (Ed.), Recent Advances in Constraints. X, 201 pages. 2003.
- Vol. 2600: S. Mendelson, A.J. Smola (Eds.), Advanced Lectures on Machine Learning. IX, 259 pages. 2003.
- Vol. 2592: R. Kowalczyk, J.P. Müller, H. Tianfield, R. Unland (Eds.), Agent Technologies, Infrastructures, Tools, and Applications for E-Services. XVII, 371 pages. 2003.
- Vol. 2586: M. Klusch, S. Bergamaschi, P. Edwards, P. Petta (Eds.), Intelligent Information Agents. VI, 275 pages. 2003.
- Vol. 2583: S. Matwin, C. Sammut (Eds.), Inductive Logic Programming. X, 351 pages. 2003.
- Vol. 2581: J.S. Sichman, F. Bousquet, P. Davidsson (Eds.), Multi-Agent-Based Simulation. X, 195 pages. 2003.
- Vol. 2577: P. Petta, R. Tolksdorf, F. Zambonelli (Eds.), Engineering Societies in the Agents World III. X, 285 pages. 2003.
- Vol. 2569: D. Karagiannis, U. Reimer (Eds.), Practical Aspects of Knowledge Management. Proceedings, 2002. XIII, 648 pages. 2002.
- Vol. 2560: S. Goronzy, Robust Adaptation to Non-Native Accents in Automatic Speech Recognition. XI, 144 pages. 2002.
- Vol. 2557: B. McKay, J. Slaney (Eds.), AI 2002: Advances in Artificial Intelligence. Proceedings, 2002. XV, 730 pages. 2002.
- Vol. 2554: M. Beetz, Plan-Based Control of Robotic Agents. XI, 191 pages. 2002.
- Vol. 2543: O. Bartenstein, U. Geske, M. Hannebauer, O. Yoshie (Eds.), Web Knowledge Management and Decision Support. X, 307 pages. 2003.
- Vol. 2541: T. Barkowsky, Mental Representation and Processing of Geographic Knowledge. X, 174 pages. 2002.
- Vol. 2533: N. Cesa-Bianchi, M. Numao, R. Reischuk (Eds.), Algorithmic Learning Theory. Proceedings, 2002. XI, 415 pages. 2002.
- Vol. 2531: J. Padget, O. Shehory, D. Parkes, N.M. Sadeh, W.E. Walsh (Eds.), Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems. XVII, 341 pages. 2002.
- Vol. 2527: F.J. Garijo, J.-C. Riquelme, M. Toro (Eds.), Advances in Artificial Intelligence - IBERAMIA 2002. Proceedings, 2002. XVIII, 955 pages. 2002.
- Vol. 2522: T. Andreassen, A. Motro, H. Christiansen, H.L. Larsen (Eds.), Flexible Query Answering Systems. Proceedings, 2002. X, 383 pages. 2002.
- Vol. 2514: M. Baaz, A. Voronkov (Eds.), Logic for Programming, Artificial Intelligence, and Reasoning. Proceedings, 2002. XIII, 465 pages. 2002.
- Vol. 2507: G. Bittencourt, G.L. Ramalho (Eds.), Advances in Artificial Intelligence. Proceedings, 2002. XIII, 417 pages. 2002.
- Vol. 2504: M.T. Escrig, F. Toledo, E. Golobardes (Eds.), Topics in Artificial Intelligence. Proceedings, 2002. XI, 427 pages. 2002.
- Vol. 2499: S.D. Richardson (Ed.), Machine Translation: From Research to Real Users. Proceedings, 2002. XXI, 254 pages. 2002.
- Vol. 2484: P. Adriaans, H. Fernau, M. van Zaanen (Eds.), Grammatical Inference: Algorithms and Applications. Proceedings, 2002. IX, 315 pages. 2002.
- Vol. 2479: M. Jarke, J. Koehler, G. Lakemeyer (Eds.), KI 2002: Advances in Artificial Intelligence. Proceedings, 2002. XIII, 327 pages. 2002.
- Vol. 2475: J.J. Alpigini, J.F. Peters, A. Skowron, N. Zhong (Eds.), Rough Sets and Current Trends in Computing. Proceedings, 2002. XV, 640 pages. 2002.
- Vol. 2473: A. Gómez-Pérez, V.R. Benjamins (Eds.), Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web. Proceedings, 2002. XI, 402 pages. 2002.
- Vol. 2466: M. Beetz, J. Hertzberg, M. Ghallab, M.E. Pollack (Eds.), Advances in Plan-Based Control of Robotic Agents. VIII, 291 pages. 2002.
- Vol. 2464: M. O'Neill, R.F.E. Sutcliffe, C. Ryan, M. Eaton, N.J.L. Griffith (Eds.), Artificial Intelligence and Cognitive Science. Proceedings, 2002. XI, 247 pages. 2002.
- Vol. 2448: P. Sojka, I. Kopecek, K. Pala (Eds.), Text, Speech and Dialogue. Proceedings, 2002. XII, 481 pages. 2002.
- Vol. 2447: D.J. Hand, N.M. Adams, R.J. Bolton (Eds.), Pattern Detection and Discovery. Proceedings, 2002. XII, 227 pages. 2002.
- Vol. 2446: M. Klusch, S. Ossowski, O. Shehory (Eds.), Cooperative Information Agents VI. Proceedings, 2002. XI, 321 pages. 2002.
- Vol. 2445: C. Anagnostopoulou, M. Ferrand, A. Smaili (Eds.), Music and Artificial Intelligence. Proceedings, 2002. VIII, 207 pages. 2002.
- Vol. 2443: D. Scott (Ed.), Artificial Intelligence: Methodology, Systems, and Applications. Proceedings, 2002. X, 279 pages. 2002.
- Vol. 2432: R. Bergmann, Experience Management. XXI, 393 pages. 2002.
- Vol. 2431: T. Elomaa, H. Mannila, H. Toivonen (Eds.), Principles of Data Mining and Knowledge Discovery. Proceedings, 2002. XIV, 514 pages. 2002.
- Vol. 2430: T. Elomaa, H. Mannila, H. Toivonen (Eds.), Machine Learning: ECML 2002. Proceedings, 2002. XIII, 532 pages. 2002.
- Vol. 2427: M. Hannebauer, Autonomous Dynamic Reconfiguration in Multi-Agent Systems. XXI, 284 pages. 2002.

Table of Contents

Invited Paper

How Formal Logic Can Fail to Be Useful for Modelling or Designing MAS	1
<i>Bruce Edmonds</i>	

Topic A: Social Theory for Agent Technology

Communicational Patterns as Basis of Organizational Structures	16
<i>Steffen Albrecht, Maren Lübcke</i>	
On How to Conduct Experimental Research with Self-Motivated Agents.....	31
<i>Luis Antunes, Helder Coelho</i>	
Cognitive Identity and Social Reflexivity of the Industrial District Firms. Going Beyond the “Complexity Effect” with Agent-Based Simulations	48
<i>Riccardo Boero, Marco Castellani, Flaminio Squazzoni</i>	
The MAS-SOC Approach to Multi-agent Based Simulation	70
<i>Rafael H. Bordini, Fabio Y. Okuyama, Denise de Oliveira, Guilherme Drehmer, Romulo C. Krafta</i>	
Organisation Modelling for the Dynamics of Complex Biological Processes	92
<i>Tibor Bosse, Catholijn M. Jonker, Jan Treur</i>	
Communication without Agents? From Agent-Oriented to Communication-Oriented Modeling	113
<i>Thomas Malsch, Christoph Schlieder</i>	
Modeling Product Awareness Rates and Market Shares	134
<i>Filippo Neri</i>	
Metanarratives and Believable Behavior of Autonomous Agents	145
<i>Mirko Petric, Inga Tomic-Koludrovic, Ivica Mitrovic</i>	
FORM – A Sociologically Founded Framework for Designing Self-Organization of Multiagent Systems	156
<i>Michael Schillo, Klaus Fischer, Bettina Fley, Michael Florian, Frank Hillebrandt, Daniela Spresny</i>	

Social Organization in a Software Agent Community with a
Non-zero-Sum Game Interaction Model 176
Matti A. Vanninen, John R. Rose

Emotion: Theoretical Investigations and Implications for
Artificial Social Aggregates 189
Christian von Scheve, Daniel Moldt

Topic B: Norms and Institutions in MAS

What Is a Normative Goal?
Towards Goal-Based Normative Agent Architectures 210
Mehdi Dastani, Leendert van der Torre

Searching for a Soulmate – Searching for Tag-Similar Partners
Evolves and Supports Specialization in Groups 228
David Hales

Norms and Their Role in a Model of Electronic Institution 240
Ioan Alfred Letia, Wamberto W. Vasconcelos

A Model of Normative Multi-agent Systems
and Dynamic Relationships 259
Fabiola López y López, Michael Luck

Integration of Generic Motivations in Social Hybrid Agents 281
Fenintsoa Andriamasinoro, Remy Courdier

Author Index 301

How Formal Logic Can Fail to Be Useful for Modelling or Designing MAS

Bruce Edmonds

Centre for Policy Modelling
Manchester Metropolitan University
<http://cfpm.org/~bruce>

“To a person with a hammer, every screw looks like a nail” (trad.)

Abstract. There is a certain style of paper which has become traditional in MAS – one where a formal logic is introduced to express some ideas, or where a logic is extended on the basis that it then covers certain particular cases, but where the logic is not actually *used* to make any substantial inferences and no application of the logic demonstrated. I argue that although these papers do follow a certain tradition, that they are not useful given the state of MAS and should, in future, be rejected as premature (just as if one had simulation but never run it). I counter the argument that theory is necessary by denying that the theory has to be so abstract. I counter the argument that logic helps communication on the simple grounds that for most people it doesn’t. I argue that the type of logic that tends to be used in these papers is inappropriate. I finish with some suggestions as to useful ways forward.

1 Introduction

During RASTA 2002 there was some discussion about the utility of formal systems for building or understanding multi-agent systems (MAS). This paper is an attempt to put my arguments. I argue that (as with any tool) one has to use formal systems appropriately. Merely following a tradition of how to use and develop a particular kind of formal system is not sufficient to ensure one is doing something useful.

In this context I wish to make it clear that I have nothing against logic. I like formal logics because they can deal with qualitative information and they can be quite expressive. However, at the end of the day¹, they are just one of a range of types formal systems that could be used – the kind of the system that is chosen is important. The point is to distinguish when and how a particular formal system is useful – this applies to formal logics as a particular case.

In short, the question is not whether to abstract from our field of study using formal systems but how. In the past, premature ‘armchair theorising’ has not helped the eventual emergence of useful theory, but rather impeded it. Formal systems (such

¹ As David Hales would say.

as logics) are not the *content* of theory but merely a *tool* for expressing and applying theory in a symbolic way – choosing the wrong kind of formal system will bias our attempts and make our task more difficult.

2 Two 20th Century Trends in Logic

Whitehead and Russell [13] showed that set theory, arithmetic and a good chunk of other mathematics could be formalised using first-order classical predicate logic. This dramatically demonstrated the expressive power of logic. Once set theory was properly logically formalised and the expressive power of set theory revealed it became clear that all mathematics could be embedded in set theory and hence be logically formalised. If any system could be shown to have an embedding in set theory, then it counted as mathematics. Thus set theory and classical first order predicate calculus was shown to *general systems*, in the sense that all known formal systems could be expressed in them (albeit with different degrees of difficulty).

In the second half of the 20th Century there was an explosion of different kinds of logic. This can be divided up into two approaches: those who were searching for the ‘one true logic’ (what I call the ‘philosophical approach’); and those who saw logic as merely a useful tool for doing complex inference (what I call the ‘pragmatic approach’). The former of these tinkered with the very structure of logic, restructuring the nature of deduction in the logic so as to attempt to match correct inference in natural language and by inventing new objects into the logic such as indices, operators, names etc. The nature of their discussions went very much by example – since they felt it was worth trying to construct the ‘one true logic’ it necessarily had to include all such cases. Logics in this vein included intuitionistic logic, free logic, relevance logic and modal logic. Due to the nature of their discussions their work tended to concentrate upon the axioms of the logic in relation to particular cases and treat the proof theory and formal semantics more as an after thought.

The pragmatic approach does not care so much about the philosophical interpretation as to what could be done with the logic. Thus, since classical first order predicate logic was generally expressive [7], they tended to work within this framework or construct simple extensions of it. For these people it was the pragmatic virtues that mattered: was it good for doing inference in; were its formal semantics checkable; was it easy to model with; and could it be used for computation (ala Prolog and its successors)? The particular logic chosen for the MAS modelling language, SDML² is a case in point – its purpose is not to capture any general theory of cognition but to provide a sound and efficient basic for the consistent firing of complex sets of interdependent rules [12].

Unfortunately the philosophical approach has tended to attract the more attention in AI. There may be many reasons for this: it may be that the association with philosophy gives it academic status; it may be that the participants truly believe that

² <http://sdml.cfpm.org>

there will be general logical systems that encode cognitive relations in ultimately simple ways; and it may be that it is relatively easy to write but difficult to criticise. Whatever the reason there has grown up a tradition in AI (and now MAS) which discusses different axiomatisations of logical systems based purely on plausibility and the ability to encode particular examples (i.e. its expressive power). It is this style of paper that I am arguing against on the grounds that, in the absence of any results, it does not merit publication.

3 Generality and Abstraction

One of the principle ways of achieving generality is to abstract away from the detail of particular cases leaving only what happens to be true of the wider domain one is considering (post hoc abstraction). Another way is to decide the structure before hand and to *choose* one's domain accordingly or else to simply ignore those aspects of those cases that seem to contradict that structure (a priori abstraction). A third way is to include a method for adapting to the particularities of each case so that the detail is preserved (adaptive generality). However it is achieved, the increased generality is obtained at a cost, a cost of lost information, relevance or computation respectively. The cost of losing information as a result of post hoc abstraction may be critical if it is the important details (w.r.t. one's goal) that are lost. The cost of restricted relevance as a result of a priori abstraction may be critical if this means that it excludes your intended object of study. The cost of increased computation may be critical if the computation is too onerous to be practical.

One well-known dynamic of philosophical discourse is that of the counter-example followed by an increase in generality: a thesis is proposed; then a case exhibited where the thesis fails; and, in response, the thesis is generalised (e.g. by adding caveats, or by being suitably elaborated). The repeated application of this process of a priori abstraction is a set of very general, but irrelevant principles. These principles may give one the illusion of relevance because the 'ghosts' of the original concepts are left as labels and symbols in the general principles and one has the impression that the relevance can be restored by the simple adding of particulars. However, if this attempted this is found to be unworkable in practice. Be clear – it is not generality or abstraction by themselves that causes this lack of relevance but the *way* the generality is achieved (i.e. a priori abstraction). Similarly – I am not arguing against generality or abstraction but that it should be done in a way that results in useful theory. Work which attempts to mimic the counter-example-generalisation process in formal logic will not result in relevant theory about MAS.

One way of clearly demonstrating that increased generality is not a sufficient reason for exhibiting a logic is that there are already many logics (and other formal systems) that are as general as possible. If a particular logic has the ability to capture a particular concept then the general one will also be able to do this. The point of inventing new formal systems is thus *entirely* pragmatic, for each system (even the general ones) will inevitably facilitate the construction of certain systems and frustrate others, just as different programming languages are good at certain tasks and

bad at others. This presence of implicit bias is not a question of the theoretical ability of the system but practical ease for us humans. This is why we neither formalise everything in set theory nor program using Turing Machines. Choosing an inappropriate formal system will bias the development of a theory in unhelpful ways, choosing an appropriate system will facilitate it [4]. Merely establishing that a particular system can express certain properties does not demonstrate that the system will facilitate a good theory, for the general systems also do this and they would (almost certainly) make formal modelling impossibly cumbersome and inference infeasible.

Thus arguing for a particular kind of formal logic on the grounds that it is able to express certain ideas, concepts or cases is very weak, for there are already formal logics that do the same (if any can). Thus, although the development of formal logics is often driven by a wish to express certain ideas, they need to be *justified* on other, stronger grounds.

4 The Need for Theory

Clearly if we are to escape simply considering individual cases and if our understanding of MAS is to inform our construction of MAS (and *vice versa*) then we will need to generalise and abstract our knowledge, i.e. use ‘theory’. The trouble is that ‘theory’ can come in a variety of levels of abstraction and a variety of forms. A natural language description is already a sort of theory because it is the result of many relevance and representational decisions – it provides a level of generalisation by facilitating the comparison of phenomena by substituting the comparison of descriptions. An MAS may be also be used as a method of producing a sort of dynamic description of a social system – this is when one attempts to program the individual agents as closely to actual accounts as possible and then check that all stages also correspond to those in the social systems at all levels of aggregation. Another MAS may be intended to represent a set of phenomena that occurs in a small set of individual cases – here the generality is restricted to a particular domain. At the other end of the scale are the ‘high theories’ of philosophy or sociology – these are ideas that are supposed to have a very great level of generality. In philosophy the theories tend to be precise but irrelevant. In contrast, in sociology the theories are relevant but often extremely difficult to pin down – they are more akin to a richly expressive language for talking and thinking about social phenomena.

I am unsure of exactly what Rosaria Conte means by ‘theory’ during her remarks during the closing panel of AAMAS 2002 (and elsewhere, see [2]). If she meant that *some* level of abstraction will be necessary for escaping from individual cases, then I agree with her – simply constructing particular MAS is not enough. However, if she is arguing that ‘high theory’ is necessary, then I disagree, for intermediate levels of abstraction also allows us to escape from single cases. For example physicists managed perfectly well to develop useful theories before the advent of their high theories, indeed they are still looking for a ‘Theory of Everything’ (TOE), even though it is clear that the situations in which such a TOE would diverge from the