

Lecture Notes in Artificial Intelligence

830

Subseries of Lecture Notes in Computer Science

Edited by J. G. Carbonell and J. Siekmann

Lecture Notes in Computer Science

Edited by G. Goos and J. Hartmanis



Cristiano Castelfranchi Eric Werner (Eds.)

Artificial Social Systems

4th European Workshop
on Modelling Autonomous Agents
in a Multi-Agent World, MAAMAW '92
S. Martino al Cimino, Italy, July 29-31, 1992
Selected Papers

Springer-Verlag

Berlin Heidelberg New York

London Paris Tokyo

Hong Kong Barcelona

Budapest

Series Editors

Jaime G. Carbonell
School of Computer Science, Carnegie Mellon University
Schenley Park, Pittsburgh, PA 15213-3890, USA

Jörg Siekmann
University of Saarland
German Research Center for Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3, D-66123 Saarbrücken, Germany

Volume Editors

Cristiano Castelfranchi
Institute of Psychology, National Research Council - CNR
Viale Marx 15, I-00137 Rome, Italy

Eric Werner
Plato GmbH, Technikzentrum Lübeck
Seelandstr. 15, D-23569 Lübeck, Germany

CR Subject Classification (1991): I.2.11, I.2, J.4

ISBN 3-540-58266-5 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-58266-5 Springer-Verlag New York Berlin Heidelberg

CIP data applied for

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 1994
Printed in Germany

Typesetting: Camera ready by author
SPIN: 10472631 45/3140-543210 - Printed on acid-free paper

Contents

Introduction

The MAAMAW Spirit and this Book <i>Cristiano Castelfranchi and Eric Werner</i>	VII
---	-----

Artificial Intelligence and Reactive Systems

Multi-Agent Simulation as a Tool for Modeling Societies: Application to Social Differentiation in Ant Colonies <i>Alexis Drogoul and Jacques Ferber</i>	3
Experiments in Multi-Agent System Dynamics <i>P.J. Kearney</i>	24
Social Aggregations in Evolving Neural Networks <i>Frederico Cecconi, Daniele Denaro, Domenico Parisi, and Ugo Piazzalunga</i>	41
An Architecture for Action, Emotion, and Social Behavior <i>Joseph Bates, A. Bryan Loyall, and W. Scott Reilly</i>	55
Game Theory Versus Multiple Agents: The Iterated Prisoner's Dilemma <i>Bjørn Lomborg</i>	69
Trust in Distributed Artificial Intelligence <i>Stephen Marsh</i>	94

Economics and Game Theory

Negotiation with Incomplete Information About Worth: Strict Versus Tolerant Mechanisms <i>Gilad Zlotkin and Jeffrey S. Rosenschein</i>	115
Equilibratory Approach to Distributed Resource Allocation: Toward Coordinated Balancing <i>Kazuhiro Kuwabara and Toru Ishida</i>	133
Strategic Interaction in Oligopolistic Markets – Experimenting with Real and Artificial Agents <i>José Castro Caldas and Helder Coelho</i>	147
Multi-Agent Coordinated Decision-Making Using Epistemic Utility Theory <i>Wynn Stirling</i>	164

Coordination and Multi-Agent Planning

The Search for Coordination: Knowledge-Guided Abstraction and Search
in a Hierarchical Behavior Space 187

*Daniel Damouth, Edmund H. Durfee, Marcus Huber,
Thomas A. Montgomery, and Sadip Sen*

Multi-Agent Planning as Search for a Consensus that Maximizes
Social Welfare 207

Eithan Ephrati and Jeffrey S. Rosenschein

Planned Team Activity 227

*David Kinny, Magnus Ljungberg, Anand Rao, Elizabeth Sonenberg,
Gil Tidhar, and Eric Werner*

DAI Tools and Applications

Some Requirements for Mobile Distributed Telecomputing Architecture 259

Carl E. Hewitt

Multi-Agent Research in the Knobotics Group 271

Yoav Shoham

A Conflict Resolution-Based Decentralized Multi-Agent Problem
Solving Model 279

Faruk Polat and H. Altay Guvenir

Decision Coordination in Production Management 295

Jacqueline Ayel

MAKILA: A Tool for the Development of Cooperative Societies 311

Karmelo Urzelai and Francisco J. Garijo

User Role in Problem Solving with Distributed Artificial Intelligent Systems 324

Lynne E. Hall, Linda Macaulay, and Greg O'Hare

Introduction

The MAAMAW spirit and this book

1 MAAMAW is recognized as the European Workshop of Distributed AI and Multi-Agent Systems. However, since its beginning, it has been characterised by a broader platform, due not to the addition of mere contour topics just to enlarge the audience (the European DAI community is not very large), but, more strategically, to the identification of some promising intersections with other AI areas or with other disciplines.

The editors of this book earnestly participated, with the other Program Committee members, in this setting, trying to give DAI a cultural role as one of the most interesting inter-disciplinary areas, and not only as a technical, specialistic sub-field of AI.

Let us mention just three of those significant openings towards other fields:

Emerging Cooperation

Since the beginning it has been recognized that for both applicative and theoretical reasons, it is relevant to develop and compare models of cooperation among agents with different levels of cognitive complexity. For a general theory of social interaction, it is necessary to develop notions like "cooperation", or "communication", or "norms, etc., both with deliberative agents and with reactive or subsymbolic agents, in which those phenomena are not intentional or planned, but just "emergent". Also in applicative contexts, such as robotics, it could be necessary to model coordination and common problem-solving among very simple agents. This perspective implies giving special attention to the problem of subsymbolic architecture and to evolutionary and dynamic modelling, so as to establish an explicit connection between DAI and that growing area called Artificial Life. That is why, since the beginning, one can find in MAAMAW models of reactive and of behaviouristic agents, approaches based on ecological concepts, the application of concepts from complex dynamic systems to the behaviour of the individual or of the group. In this book, we find the reactive approach of Ferber's group to the simulation of a population of ants; neural-net organisms in evolution (Parisi et alii); "potential fields" that control the complex dynamic behaviour of simple agents (Kearney).

Logic of Action and Agent Architecture

MAAMAW is interested not only in modelling coordination, negotiation, and so on, but also in connecting such social phenomena and deriving them from appropriate architectures of the social agent, or vice versa, in showing how society and groups shape the minds and the behaviours of their members. The former alternative depends on the adoption of a "methodological individualism" or of a so called "psychological" perspective, while the latter alternative depends on a so called "sociological foundation" of DAI (Gasser) which claims that Society, rather than the individual, is the prius: "Society comes first". Thus, a dialogue between, on one side, current studies on the BDI (Belief-Desire-Goal) architectures and in general the theory of action, intention, belief, and, on the other side, Multi-Agent Systems, is necessary. Which agent's goals or beliefs its participation in social co-operation? Is it benevolent or self-interested? Is it trusting the other agents or evaluating their sincerity, competence and reliability? Is it autonomous or slavish? How a collective mind is formed? How are joint intentions related to the intentions of the individuals? In the book one can find papers relative to that important trend of convergence between DAI and agent/action theory and architecture (for ex. Shoham; Kinny et alii).

Social modelling and Social Simulation

Like in Cognitive Modelling, so in Social Modelling AI is interested in describing and explaining (modelling) real social phenomena (such as negotiation, persuasion, alliances, conflicts, social hierarchies, etc.). This is because, on one side, contributing to a scientific theory of social behaviour and social structure could be one of the aims of AI (like that of contributing to Cognitive Science and to the understanding of human mind); on the other side, Artificial Intelligence systems should be inserted in and support the work and the communication among humans, both in organizations and in private interactions. Therefore, understanding, modeling and dealing with human interaction are necessary conditions for an effective integration of intelligent technologies in a society.

Besides, many DAI models and platforms elicit issues which are very close to those related to the computer simulation approach in Social Sciences. That growing area has many contacts with this perspective on DAI. These contacts were explicit in MAAMAW: not only were the MAAMAW'92 chairs also co-organizers with Nigel Gilbert and Jim Doran, of the international workshops on "Simulating Societies" in 1992 and in 1993, but also, as you can see from this book, this was one of the topics of interest of our workshop.

Now, we think, the motivation for the title of the book: *"Artificial Social Systems"*, becomes clear. This is not, of course, the proposal of a new discipline (like AI), but just the stress on an important convergence among different domains: formal studies in Social Sciences, computer simulation in sociology and economy, social phenomena in Artificial Life, modelling interaction and organization in Artificial Intelligence, the creation and support of virtual communities of cooperating hybrid agents (individuals, data bases, knowledge based systems, intelligent software agents, robots, organizations). All this requires a theoretical understanding, an operational or formal characterization, the computer modelling of social interaction, and the creation of some "Artificial Social System".

This book is not reducible merely to the Proceedings of the MAAMAW workshop held in S.Martino al Cimino (Italy) in 1992. For this reason as well it has a title that, on one side, stresses the special focus of the '92 workshop, and on the other side is oriented towards students who come not only from the domain of Distributed Artificial Intelligence, but also from general AI and Computer Science, and from Social Sciences. We asked the authors to give us an extended and updated version of their papers, and in some cases the works have turned out to be of a higher quality and different from the preliminary version presented at the workshop. We asked the invited speakers to contribute to the book with some current work of theirs. Therefore the research presented in the book is current and original..

2 The round-table discussions on "AI vs DAI" (Chair Van Dyke Parunak), and on "(D)AI, Social Simulation, and Social Sciences" (Chair: John Campbell) were one of the most interesting aspect of MAAMAW'92 . Let us propose the main questions of these discussions to an enlarged audience.

AI versus DAI: Where are we ? Where are we going ?

1) What is the relationship between DAI and classical AI ? Does DAI represent a change of paradigm within AI, i.e., a different philosophy of mind, an alternative view (social-distributed vs individual) of intelligence (of action, knowledge, meaning, etc.)? Or, at least, does DAI represent some salient aspect of a shift of paradigm in AI, characterized by more context-dependent and adaptive agents, by more changeable and unpredictable environments, by interaction as a necessary means for achieving desired outcomes?

Even though DAI might not be seen as a new paradigm, what is its contribution to the general problems of AI? Which are the advantages offered by "distributed" systems?

How have the biases of single-agent AI influenced the fields of AI and DAI ? Are there new biases (say, collectivistic ones) introduced by DAI ? In general, what are in your view the limits of the theoretical and methodological broad lines of both AI and DAI ?

What would be a "Turing Test" for DAI ?

What is a critical real-world problem whose solution by a DAI researcher would most impress upon the general AI community the importance of DAI ?

2) Is DAI concerned with interactions among intelligent agents, or are also social phenomena among reactive or sub-symbolic agents, as well as forms of emergent cooperation, within its province? What are the relations and the differences between DAI and distributed processing systems? In other words, what is an "agent" and how is it different from a "processing unit" or from a "node" of a network, or from an "actor"?

3) Is the exclusive or main aim of DAI that of designing, on the grounds of formal theories and experiments, systems (methods as well as artifacts) effective for problem solving through interaction? Or, is it also (or mainly) that of describing (however in idealized terms) "natural" social phenomena, and of providing some theoretical or experimental explanation of such phenomena?

In other words, does DAI propose theories and models of natural social conditions, or does it have a "prescriptive" purpose (such as operational research or the theory of rational decision) with regard to the methods for social interaction and efficient collective action?

4) Is DAI a relatively unified sub-discipline of AI? Is it useful to view "distribution" as a key problem that ties together the work of various researchers, using various methodologies?

Is it useful a better theoretical unification of DAI? and where is it to be found? Many crucial problems and concepts have been identified (such as coordination, organization, commitment, negotiation, reasoning about others' minds) but what is lacking is a unifying framework. What is the role to be played in this framework either by the models of individual mind and action, or by the models provided by social and management sciences, or by formal theories? Do folk notions of interaction, based on concepts like coordination, cooperation, and negotiation, exert too much influence on DAI research? Not enough?

5) How can the gap between implemented DAI systems and DAI theories be bridged? Should we demand that theoretical DAI have an obvious impact on implemented systems? What computational expectations should we have of DAI theories?

(with the collaboration of Jeff Rosenschein)

DAI, Social Simulation and Social Sciences

As for the panel on "(D)AI, Social Simulation and Social Sciences", three points emerged from the debate:

The general interest of DAI in social sciences models and problems, and vice versa;

The notion of "emergence" and the problem of reductionism;

The importance of more adequate, non metaphorical social notions in DAI, such as "groups", "negotiation", "norms", "power", "role", etc.

The notion of Emergence

As for the notion of Emergence, the discussion has highlighted the presence of a confusion among different kinds or notions of "emergence", some implying the temporal or developmental dimension, other implying a relationship between subsymbolic and symbolic, others relative to the micro-macro link in sociology (i.e. the link between the individual features, actions, and intentions, and the establishment of a collective mind or of a collective unintended functional behaviour). It was also argued against the reduction of emergent functional cooperation to subsymbolic or reactive agents, as if it were impossible to have functional cooperation among deliberative agents beyond their awareness and intentions.

Impact on Social Sciences

DAI is expected to have, on social sciences, (Sociology, Social Psychology, Economy, Pragmatics, Organization Science, etc.) an impact comparable to the revolutionary impact of AI on Psychology (with the rising of Cognitive Science). Many social sciences need, on one side, well defined, formal or operationalized, conceptual instruments; on the other side, instruments for experimenting on mind and on social forms and evolution.

DAI could significantly contribute to both these theoretical and methodological aspects. It can furnish for example, formal or computational models of social-cognitive agent architecture; it can provide platforms for experimenting social coordination, organization, or change.

In particular, we think that, with respect to some important formal models of social action (like Game Theory), or with respect to some more traditional, mathematical approaches to Social Simulation, DAI uniquely allows to predict or discover, by simulation, social complex behaviour or trends emerging from a complex but modelled cognitive activity. In general, DAI is unique in accounting jointly for the social dynamics and for the internal architecture and dynamics of the agent (such as reactivity, or planning, or belief revision, etc.).

From Social Science to DAI

The other direction of the relation is important as well: the one from Social Science to DAI. Social Sciences could offer to DAI both:

Many important, puzzling, and well studied but unresolved problems of social interaction (such as reciprocity and the possibility of cooperation among self-interested agents; the source of norms and authority; or the relationships between efficiency/efficacy and the kind of group and organization; and

Some relevant formal approaches and theories about human action (at least in very idealized conditions), like rational decision theory, economic theory, Game theory. AI should use and absorb these theories critically, without forgetting that they are under debate, from both a philosophical and an empirical point of view, in their original disciplines.

Metaphorical notions

As for this last point, frequently enough DAI social notions are merely metaphorical: they just remind the correspondent social notions (like contract, negotiation, commitment, team, etc.), but do not account for many crucial aspects of these notions in analysing true social phenomena. We think that social notions in DAI should be more descriptively adequate, closer to (but more precise than) their corresponding sociological notions. This is not only for scientific adequacy, but also because there is the need in application as well, for a subtler discrimination and more articulated conditions. For example, we must be able to discriminate among different forms of "cooperation", without mixing up social exchange and altruistic help, forced obedience and spontaneous collaboration.

The book, of course, not only refers to the topics mentioned above, but is also rich of significant DAI and Multi-Agent languages, models, and platforms. Also some important problems of application are discussed. It is worth reminding that MAAMAW organizes each year an Olympics for some demonstrations of DAI applications.

Let us now shortly describe the parts and the chapters of the book.

3 The book is divided in five sections: Artificial Life and Reactive Systems, Social Modelling and Simulation, Economics and Game Theory, Multi-Agent Planning, DAI Tools and Applications

Artificial Life and Reactive Systems

Alexis Drogoul and Jacques Ferber in their paper describe a general model of simulation of complex societies based on the simulation of the behaviour of its individuals. The aim is to contribute to the understanding of "emergence" in ecological and sociological systems; to enlighten the mechanisms of sociogenesis and the micro-macro link. They also present a multi-agent simulation system EMF based on the definition of reactive agents whose behaviour is governed by the selection of simple competing tasks due to stimulus perception. On such a basis they describe the MANTA project which is aimed at modelling an entire ant society and its evolution. The architecture of the agents, their various roles and behaviours are specified, and some early experimental results are presented.

Paul J. Kearny applies field theory to the behaviour of individual agents. To determine such a behaviour he uses a mechanism analogous to potentials that determine the motion of particles in a potential field. Applied to the agents potentials can be thought of as representing their motivations (fear, hunger, etc.). In such a sub-symbolic framework (with reference to a certain MAAMAW tradition), a series of experiments investigate the dynamic behaviour of systems of coupled agents. The behaviour of the coupled systems can be extremely complex even when agents are very simple. Insightful considerations arise concerning animal instincts, collective intelligence, and emerging cooperation

Domenico Parisi, Ugo Piazzalunga, Federico Cecconi, and Daniele Denaro present simulations of populations of simple organisms living together in the same environment, in order to study social aggregation as an emerging phenomenon. The evolution of neural networks is simulated using genetics algorithms. Spatial aggregations emerge evolutionarily, not only as a by-product of the spatial distribution of resources, but also as an advantageous adaptation of living inside social groups ("information centres"), and as a pre-condition for learning from others.

Joseph Bates, A. Bryan Loyall, W. Scott Reilly present the emotions and social behaviour of Lyotard, a simulated house cat in an Oz micro-world. The Oz project at Carnegie Mellon University is developing technology for artistically interesting simulated worlds. An agent architecture, called Tok, is described, which support perception, reactivity, goals, emotions, and interaction. Behaviours are related to emotions and goals. Even a quite simple agent like Lyotard exhibits interesting attitudes and social behaviours.

Social Modelling and Simulation

Bjorn Lomborg's work faces with the evolution of cooperation and of its stability. This research integrates noise, misunderstandings (that are fatal for the well known Tit-for-Tat strategy) in a real evolutionary framework with a large pool of competing strategies, and with thousands of individuals and generations. The simulation shows that cooperation stability is possible even under high levels of noise. The multi-agent approach and the computer simulation offer interesting solutions to this problem of cooperation ensuing in the Iterated Prisoner's Dilemma, that analytical Game Theory cannot be expected to solve. The main result

is the fact that the unexpected stability of cooperation is due to the coexistence of many different behavioural strategies that compose the emergent phenomenon. In political terms: pluralism seems to be a basis for the ensuing and stability of cooperation.

Stephen Marsh analyses the role of Trust from the point of view of an agent interacting in a Multi-Agent world in which there could be irresponsible and malevolent agents. He defines a specific notion of Trust which capture some important features of the common sense notion, but not all its connotations. Thanks to these features this notion exhibits its usefulness in a DAI context allowing the agent to take some risks in social cooperation, but also to be prepared to cheating and to minimising potential damage. We think that this kind of notions (like the related notions of "reputation", "responsibility", "reliability", "credibility", etc.) will be crucial in the development of both the theory and the applications, insofar as to treat with truly autonomous, selfish, and heterogeneous agents will be needed.

Economics and Game Theory

MAAMAW anticipated the currently growing area that at IJCAI'93 has received the name of "Artificial Economics". Also in this volume one finds papers in which either DAI concepts and architectures are used to model economic interactions and phenomena, or economic constructs and theories are used in AI to account for multi-agent decisions or dynamic equilibrium.

A special significance is attached to Game Theory, given its role for a formal approach to many social phenomena: from economics, to politics, to ethics, to sociology. Game Theory, in a certain sense, is playing the role of an "interface" between AI and Social Science (see also Lomborg, and Marsh). This is acceptable, provided that one does not assume Game Theory just as a technical instrument, but as a complex approach with many philosophical and methodological problems, and with many limitations as a general theory of social behaviour.

In the paper by *Gilad Zlotkin* and *Jeffrey Rosenschein* we are dealing with the problem of the efficiency and the stability of outcomes in a negotiation process among utilitarian agents. These agents have incomplete information about each other. In particular, they know about the goals of the other agents, but they ignore the value, the "worth" that the agent attach to their goals. To negotiate, an agent must declare this worth; thus, which is the best declaration strategy the agent should adopt for increasing its utility?

Kazuhiro Kuvabara and *Toru Ishida* discuss the hypothesis whether the market metaphor is the appropriate foundation for distributed systems. They analyse two classical microeconomic approaches to resource allocation problems: the explicitly *cooperative* approach, in which each agent knows the global utility function, calculates the marginal utility, and exchanges its value with the other agents; the *competitive* approach, in which agents compete with one other, yet they are expected to contribute to global utility through an *invisible hand*. They propose a third approach (*symbiotic*) that neither involves explicit cooperation nor competition among the agents. They show through simulation that the proposed symbiotic approach is effective for a distributed resource allocation where a global objective function should be guaranteed, and they claim that this approach reduce communication between agents.

The paper by *Jose' Castro Caldas* and *Helder Coelho* proposes DAI techniques as valid instruments to operationalize and make experiments about economic theories. DAI can provide a computational test-bench for research in economic interaction. In particular, in the paper the DAI agents play the role of producers in an oligopolistic market, and they exhibit behaviour similar to human agents playing an analogous role.

Wynn Stirling's paper presents an agent architecture based on the theory of the "epistemic utility" of knowledge, that accounts for agent's preferences, beliefs, decisions, risk taking. This model is applied to a multi-agent situation to produce coordinated decision making. Coordination is obtained among heterogeneous and autonomous agents thanks to the fact that agents are able to reason about the epistemic utility of the other agents.

Coordination and Multi-Agent Planning

In the paper by *Edmund H. Durfee, Daniel Damouth, Marcus Huber, Thomas Montgomery, and Sandip Sen*, a very interesting and powerful analysis of coordination as a search process is presented. The individual agents must find appropriate activities that allow them to achieve individual and collective goals. It is highlighted how traditionally distinct coordination techniques can be viewed as search processes. Their similarity, but at different levels of abstraction, is explained. Temporal and interdependence aspects of collective activities are accounted for, so that search coincides with organizational design or distributed resource scheduling. The paper stresses the role of abstraction and knowledge in this distributed search process. Promising applications are presented about distributed meeting scheduling, and coordinating multi-robot arms.

To reach consensus in their coordination, agents could negotiate, but they could also vote. Which voting mechanisms prevent the agents from manipulation by untruthful agents? In *Etan Ephrati* and *Jeff Rosenschein's* paper, this vote mechanism is applied to group planning. Just expressing their local preferences step by step, the agents incrementally construct a common plan that brings the group to a state maximising the global utility. It is clear the interest of this new multi-agent planning technique in assuring social welfare among heterogeneous agents, and in repressing insincerity.

The work by *David Kinny, Magnus Ljungberg, Anand Rao, Elizabeth Sonenberg, Gil Tidhar, Eric Werner* provides a framework for planned team activity in which to explore such aspects of joint action as: team formation, role assignment, joint plan execution, and co-ordinated recovery from failure. A language for specifying joint plans at a team level is introduced. It is assumed that plans are supplied in advance to agents rather than being generated as required. Intentions capture agents' commitments to joint activity. Joint intention is analysed as a conjunction of individual intentions together with mutual beliefs about the intentions of other individuals. Joint intention takes into account also the responsibility of members of a team to communicate their failures to other members. The formalism provides an effective framework for reasoning about joint actions, roles, skills, teams, commitments, and co-ordination in executing complex plans.

DAI Tools and Applications

Carl Hewitt's paper introduces components (*Participants, Communications, Events, Relationships, Summaries*) and services of *Joint Activities* in Mobile Distributed Telecomputing Architecture. Services of Presentation/Interaction management, Meeting/Encounter management, Discourse management, Organizational Relationship management, Project/Task management, and Process management, are identified. In particular, it is shown how *Electronic Organization* and *Joint Activities* can improve the rapidity of response and the robustness of the distributed cooperative system.

Yoav Shoham provides a summary of recent research with the Agent Oriented

Programming (AOP), and an overview of this well known approach. AOP can be viewed as a specialization of object-oriented programming, where the state of the agent is its *mental state*. In fact the state of an agent consists of beliefs, choices, capabilities, commitments, etc. An extension of standard epistemic logic is introduced to capture those *mental states*. It includes temporalized knowledge and belief operators, and operators for capability, choice, and commitment. The *agent programs*, which control the agent behaviours, include also some speech-act primitives, like *informing*, *requesting*, *offering*, for communicating with other agents.

One of the important problems faced in a cooperating community of experts is how to detect and resolve conflicts occurring at any phase of problem solving. *Faruk Polat* and *Altay Guvenir* present a model in which the agents are not assumed to have a global conflict resolution knowledge. Each agent has its own conflict knowledge which is separated from its domain level knowledge. The conflict resolution knowledge is not accessible and known by others. Each agent involved in a conflict is free to choose a resolution scheme according to its self-interest. This decentralized problem-solver model is described by using an example in the domain of collaborative office design.

Jacqueline Ayel's paper analyses production management decisions in Computer Integrated Manufacturing. The coordination of decisions cannot be viewed as a mere problem of Data Sharing between decision maker systems: the management of conflicts between them should be taken into account. The paper presents an architecture in which the "supervision" (coordination and synchronization) is analysed as a distributed task. A layer of cooperating "Local-Controllers" is introduced. Each Local-Controller is a knowledge-based reactive system which uses a blackboard mechanism. Various kinds of cooperation and various decisional politics are presented.

MAKILA is a tool, developed by *Karmelo Urzelai* and *Francisco Garijo*, for the construction of societies of distributed agents, that use negotiation to cooperate. The model is supported by a blackboard architecture that is kept hidden, under the user's interface functions. A predefined agent traces every established contract and guarantees a social control. The model is tested on an Urgent Medical Assistance problem. The main advantages of the system are from the knowledge engineering point of view.

Lynne Hall, *Linda Macaulay* and *Greg O'Hare's* paper focuses on how to provide an ideal interaction situation for the users in relation to a DAI system. They look at User-DAI system interaction as at a Group problem solving and analyse the nature of the problem-solving environment, of group interaction (antagonist, cooperative, etc.), of the activity. The forms of interaction are clarified through consideration of the CIDIM application within the ARCHON project. Various possible user roles are identified. The conclusion (premise to a methodology for DAI system and user interface design) is that for cognitively complex tasks, the optimum is when the user is interacting with the DAI system as a partially integrated entity.

Cristiano Castelfranchi (*) and Eric Werner

Roma, April 1994

(*) MAAMAW'92 and this work were partially supported under the "Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo" of the Italian National Research Council

The chairs would like to thank Amedeo Cesta and Maria Miceli, and the Institute of Psychology of the Italian National Research Council, for the wonderful success of that workshop. Scientific quality was also guaranteed by the very collaborative work of all the PC composed by:

Magnus Boman (Stockholm University and R.I.T. - Sweden)
 John Campbell (University College, London - United Kingdom)
 Rosaria Conte (IP-CNR, Roma - Italy)
 Yves Demazeau (LIFIA/IMAG, Grenoble - France)
 Mauro Di Manzo (Universita' di Genova - Italy)
 Jacques Ferber (LAFORIA, Paris - France)
 Julia Galliers (University of Cambridge - United Kingdom)
 Graca Gaspar (Facultade de Ciencias, Lisboa - Portugal)
 Heikki Hammainen (Nokia Research Center, Espoo - Finland)
 Hans Haugeneder (Siemens AG, Muenchen - Germany)
 George Kiss (The Open University, Milton Keynes - United Kingdom)
 Paul Levi (Technischen Universitaet, Muenchen - Germany)
 Frank v. Martial (DETECON, Bonn - Germany)
 Jean-Pierre Muller (Universite' de Neuchatel - Switzerland)
 Eugenio Oliveira (Universidade de Porto - Portugal)
 John Perram (Odense Universitet - Denmark)
 Jeffrey Rosenschein (Hebrew University, Jerusalem - Israel)
 Walter Van de Velde (Vrije Universiteit, Brussels - Belgium)
 Peter Wavish (Philips Research Lab, Redhill - United Kingdom)
 Gilad Zlotkin (Hebrew University, Jerusalem - Israel)