科技资料

Decentralized A.I.

Decentralized A.I.

Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World
Cambridge, England, August 16–18, 1989

Edited by

YVES DEMAZEAU

Chargé de Recherche au Centre National
de la Recherche Scientifique
Laboratoire d'Informatique Fondamentale
et d'Intelligence Artificielle
Grenoble, France

and

JEAN PIERRE MÜLLER

Professeur à l'Université de Neuchâtel
Institut de Mathématiques et d'Informatique, IMI
Neuchâtel, Switzerland

1990

NORTH-HOLLAND
AMSTERDAM • NEW YORK • OXFORD • TOKYO

ELSEVIER SCIENCE PUBLISHERS B.V.
Sara Burgerhartstraat 25
P.O. Box 211, 1000 AE Amsterdam, The Netherlands

Distributors for the United States and Canada:
ELSEVIER SCIENCE PUBLISHING COMPANY, INC.
655 Avenue of the Americas
New York, N.Y. 10010, U.S.A.

Library of Congress Cataloging-in-Publication Data

European Workshop on Modelling Autonomous Agents in a Multi-Agent World (1st : 1989 : Cambridge, England)
Decentralized A.I. : proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Cambridge, England, August 16-18, 1989 / edited by Yves Demazeau and Jean-Pierre Müller.
p. cm.
Includes bibliographical references.
ISBN 0-444-88705-9
1. Artificial intelligence--Congresses. I. Demazeau, Yves. II. Müller, Jean-Pierre. III. Title.
QC334.E99 1989
006.3--dc20 90-39141
CIP

ISBN: 0 444 88705 9

Printed in the Netherlands

Foreword

The papers in this book have been first presented at the 1st European Workshop on Modelling Autonomous Agents in a Multi-Agent World, held in King's College at Cambridge, England, in August 1989. The format of this book is based on the collection and categorization of these papers that have been revised by the authors after the workshop. The reasons that guided the choice of both the title and the structure of this monograph are explained in the first paper.

The idea for the workshop originated in the COST-13 MAGMA Project (Modelling an Autonomous aGent in a Multi-Agent world) which contributed to recent developments in Decentralized Artificial Intelligence. The project took place in the framework of the European COST-13 Action (Commission of the European Communities, DG XIII). We are grateful to Jean-Claude Latombe who inspired this project and we would like to thank the Commission for supporting it.

We would like to thank John Campbell for organizing the workshop, the Commision of the European Communities for both contributing to the genesis of this workshop and supporting it, and every participant in the COST-13 MAGMA project for their continuous and valued help.

We are grateful for the support of our institutions, the "Centre National de la Recherche Scientifique", the "Laboratoire d'Informatique Fondamentale et d'Intelligence Artificielle" in Grenoble, the "Université de Neuchâtel", and the "Institut d'Informatique et de Mathématiques", for providing us with an excellent research environment for many years.

The first editor also thanks the Artificial Intelligence Laboratory at the Vrije Universiteit Brussels and the SWIFT company for providing the environment and support for Decentralized Artificial Intelligence which made this compilation possible.

Yves Demazeau
Jean-Pierre Müller

Contents

Interactions between Reasoning Capabilities

Interaction between Perceiving Capabilities

Practical Issues for Information Exchange

Introduction

DECENTRALIZED A.I.
Yves DEMAZEAU & Jean-Pierre MÜLLER (Eds.)

DECENTRALIZED ARTIFICIAL INTELLIGENCE

Yves Demazeau[1]

LIFIA-IMAG/CNRS
46, av. Félix Viallet,
F-38031 Grenoble Cédex
FRANCE
demazeau@lifia.imag.fr

Jean-Pierre Müller

Université de Neuchâtel
Chantemerle, 20
CH-2000 Neuchâtel
SWITZERLAND
inmuller@cnedcu51.bitnet

Decentralized Artificial Intelligence (DzAI) is concerned with *the activity of an autonomous agent in a multi-agent world.* Since DzAI focuses on modelling an agent we introduce a very simplified model of an agent to highlight some essential features of an autonomous agent in a multi-agent world. The consideration of both the *locality of a task* and the *capabilities of an agent* to achieve a task, allows us to distinguish four *kinds of behaviour* : cohabitation, cooperation, collaboration, and distribution. The consideration of the *kinds of information that can be exchanged* between several agents allows a fine-grained description of the *interactions between agents.* Such considerations, by emphasizing the agent modelling, offer an interesting and systematic framework for studying multi-agent problems [2] .

1 Scope and Goal

This paper is an attempt to structure the activity in multi-agent modelling by an approach that is different than classical Distributed Artificial Intelligence (DAI) : Decentralized Artificial Intelligence (DzAI). We first introduce a definition of these two points of view and compare them.

1.1 Distributed versus Decentralized Artificial Intelligence

DAI, as defined in [13], is concerned with the *collaborative solution of global problems by a distributed group of entities.* The entities may range from simple processing elements to complex entities exhibiting rational behaviour. The problem solving or task performing is *collaborative* in the sense that mutual sharing of information is necessary to allow the group as a whole to produce a solution, or to successfully accomplish a *global* task. The group of entities is distributed in that both control and data are logically, and often geographically, distributed.

[1]The author is currently at the Vrije Universiteit Brussels. Artificial Intelligence Laboratory, Pleinlaan 2, B-1050 Brussels, demazeau@arti.vub.ac.be, with the support of CNRS, and having a SWIFT AI chair.

[2]This work has been done in the framework of the European CEC COST-13 Project MAGMA "Modeling an Autonomous Agent in a Multi-Agent World".

DzAI is concerned with the *activity of an autonomous agent in a multi-agent world.* The word "agent" is used in a broad sense to designate an intelligent entity, acting rationally and intentionally with respect to its own goals and the current state of its knowledge. By "autonomous" agents, we mean that each agent has its own existence, which is not justified by the existence of other agents. We consider several autonomous intelligent agents which coexist and may collaborate with other agents in a common world. Each agent may accomplish its own tasks, or cooperate with other agents to perform a personal or a global task.

In the light of these definitions, DAI and DzAI have a common interest in the behaviour of distributed entities. However, in DAI, a global task is initially defined and the problem is then to design the distributed entities to enable the performance of this global task. The main issue is to study the distribution and the collaborative resolution of the given task. In DzAI, the decentralized autonomous entities are initially defined. We then study how these entities are able to achieve tasks that may be personal or may interest several entities. The main issue is to study the structure of the autonomous entities to provide insights into what kinds of problems these entities are able to solve or what tasks they are capable of performing.

1.2 Motivations for DzAI

Most research works in AI have considered a single agent having complete control over the world. To design one agent having full capability to control its environment appears a difficult task.

- the agent has to deal with multiple, uncertain, contradictory sources of information
- the agent has to deal with multiple, contextual, conflicting goals
- the agent has to map its goals into its restricted perception and acting capabilities

To consider a multi-agent world is an improvement from the point of view of both a single agent and a society of agents :

Designing more powerful and more autonomous agents : through developing new planning or solving methods that process incomplete and uncertain knowledge by communication with other agents, DzAI can provide insights into improving and increasing the reasoning and decision capabilities of each agent for the achievement of its own goals. The maintenance of incompatible information or goals, and the resolution of the incompatibilities when necessary, occur naturally in a multi-agent framework.

Understanding social behaviour : through collaboration between autonomous agents, DzAI can not only improve the intrinsic power of the capabilities of each agent but also model social behaviour emerging from the agent interactions. DzAI can provide keys for understanding the interactions among intelligent beings who belong to various groups or societies, since it is based on the very structure of the interacting entities.

Since in the DzAI approach the agents are defined first, we begin by introducing a very simplified model of an agent. Based on this description, we propose some criteria on which the multi-agent problems can be classified, and use them to describe different papers in this field.

2 A generic agent

Our intention in describing a generic agent is to define the minimal features from which we can explain the various interactions in a multi-agent world. The purpose of the following description is not to specify what the architecture of an agent is or should be but to illustrate in which terms its behaviour can be described.

Any agent can be thought of as having representations of the world or of the problem it has to solve. Let us just call these representations *knowledge* whatever form, implicit or explicit, they may take. This knowledge can be given or acquired through perception or communication with other agents. A good discussion of the nature of knowledge can be found in [20].

It is always possible to abstract *goals* from the observation of the behaviour of agents. These goals do not need to exist explicitly inside the agent. The goals can be encoded in the algorithms or explicitly given or acquired by the context or by communication with other agents.

From knowledge and goals, an agent can be thought of as having to consider a set of possible solutions or plans to acheive its goals. For the sake of generality, let us call it the set of *possible solutions*. An agent does not need to be able to derive all the possible solutions but only a part of these depending on its *reasoning capabilities*. This last notion allows us to take into account limited reasoning [14,7].
When various possible solutions (or plans) are potentially applicable, a decision must be made among them to choose the best one from the point of view of the agent. Let us call it the *choice* of the agent. For planning agents, this choice is called the intention of the agent. Due to agent limitations, the choice can be wrong from an external point of view. The limitations can be taken into account by introducing a model of the *decision capabilities* of the agent.

This proposal is illustrated by the following figure 1.
In the following paragraphs, we detail some features of the autonomous agent in a multi-agent world and show how this generic agent can be used to explain interactions between more or less autonomous agents.

3 Some features of the autonomous agent

3.1 Dynamics of the world

Each agent has at best incomplete, uncertain and partly erroneous knowledge about the world and about other agents. As a consequence, it is not possible for an agent to safely plan actions or to find solutions that are applicable in the long term. Un-

Figure 1: Agent model

certainty about the actual choices of other agents requires that an agent carefully monitors plan execution, interweaves plan generation and execution, and frequently updates its knowledge, goals, and choices. In the case of DzAI, the models used to update knowledge about the changing environment will be expressed in terms of agent behaviour.

3.2 Kinds of agent behaviour

Each agent may anticipate the actions to be executed by other agents, and may interact with them in order to satisfy its goal with respect to the choices of the other agents. Our presentation of the several kinds of agent behaviour is based on the two following criteria :

the locality of the task to be performed by the agent : *personal* (local) or *interpersonal* (global)

the capability of the agent to perform the task alone : *able* or *unable*

According to these criteria, there are naturally four ways to describe how to take into account the other agents evolving in the same world :

Cohabitation : the agent has to successfully accomplish a task and it is able to execute it alone. This point of view is equivalent to classical AI problems of planning in a mono- agent environment, but whose dynamic behaviour can be explained in terms of actions of similar entities with which conflicts may occur.

Cooperation : in order to perform a personal task, an agent will have to cooperate with others either because it is not able to accomplish it itself (restricted possible solutions), or because others can successfully accomplish it more efficiently (e.g. within a shorter interval of time). This is something deductible by the decision process of the agent.

Collaboration : some global goals may concern all agents and may be realized individually by several agents. The main problem deals with electing one of the agents to carry out the task. The redundancy of the capabilities of these agents contributes to the robustness of the whole system.

Distribution : finally, some global goals can only be achieved by several agents collectively. The main problem deals with splitting the global task and distributing it to the cooperative agents. This kind of behaviour, increasing the effectiveness of the entire system, can be observed in DAI systems.

Obviously, other dimensions of the multi-agent problems can be taken to classify the kind of agent behaviour. As an example, Werner [23] describes various styles of cooperation between agents ranging from the totally cooperative and selfless to the absolutely selfish.

3.3 Kinds of information exchange

One way for exchanging information with other agents is to communicate with them. An agent may also perceive the world through sensors such as visual sensors. Both the perception and the communication acts may be considered as actions to be planned. Through acquiring knowledge about other agents, an agent may be able to recognize the actions, and eventually will be able to infer the choices of, other agents (plan recognition). We can distinguish three types of exchange (through perception or communication) between agents evolving in a common world, according to the type of information that is effectively exchanged.

Knowledge : due to differences and due to the incompleteness in perceiving the environment, two agents may have either complementary or conflicting descriptions of a shared situation. In the first case, reasoning about the knowledge of other agents can be useful for supporting exchange of information as pointed out in [1]. Conflicts arise in the case where contradictory descriptions are present ; the problem is to negotiate in order to decide which description better matches reality. This last problem has been pointed out in [9], and its resolution involves exchanges of justifications.

Possible solution : the exchange of possible solutions arises when two or more agents have to agree on a common solution or plan. The process can be thought of as finding the intersection of the possible solutions issued by each agent. Due to different reasoning capabilities, this intersection can be empty ; in this case, reasoning about the reasoning capabilities of other agents may be useful to deduce which knowledge must be communicated to produce other possible solutions. This kind of protocol has been suggested by [19].

Choice : given a set of common possible solutions, one has to agree on a common choice when cooperation is needed. One possibility is to avoid the problem by agreeing on choosing the first solution uttered in the intersection of possible solutions. A more elaborate but passive solution requires reasoning about the decision capabilities of each agent. For example, a possible solution which is very highly rated by an agent may be rejected because the conjoined income from the point of view of both agents is in fact less interesting than the single income of each solution [10]. The communication of choice also occurs when an agent requests another agent to do something. However, when a choice is not accepted by another agent, negotiation must occur [3].

4 Interaction between several agents

4.1 Strong interaction between decision capabilities

Interaction between several agents cannot only be described by single message exchange. Usually, communication occurs through sophisticated protocols for informing, requesting or convincing. For example, requesting is not only composed of an order to do something but is usually followed by an acknowledgement of the acceptance of the request or even a set of exchanges to agree upon the intended meaning of the request. These protocols can be described from an external point of view as in [2].

They also can be explained in terms of conflicts between what the agent wants to do or more specifically about what the agent considers as most important. This last point requests reasoning about the decision process of each other [10] and about the influence each agent can have on the other (social power) [3]). A protocol for decision making in distributing tasks among several agents based on the contract net paradigm is described in [22].

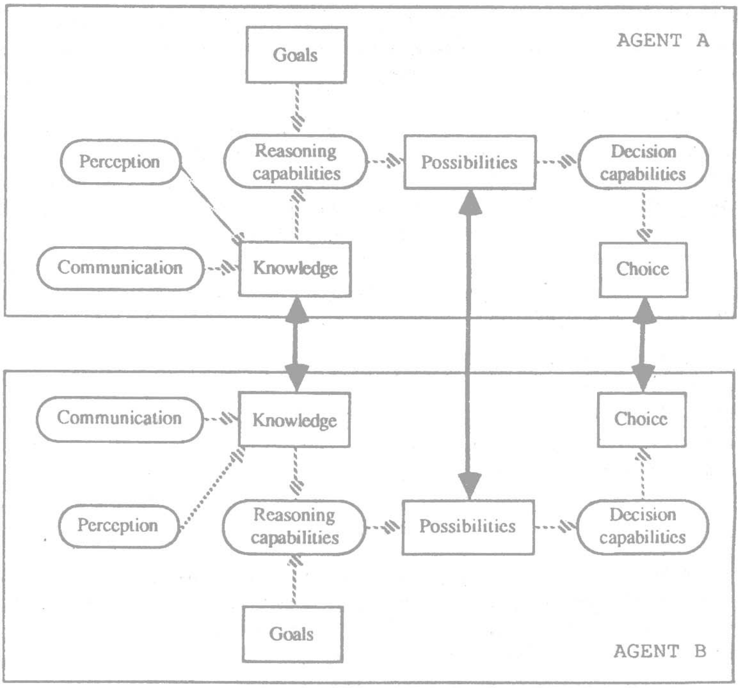

Figure 2: Strong interaction between decision capabilities

4.2 Medium interaction between reasoning capabilities

A weaker interaction than a full protocol for influencing each other occurs when the agents want to simply take into account what the other agents intend to do. One example is the simplified protocol where agents try to produce multi-agent plans (plans with actions from both agents) by using the model they have of their own reasoning capabilities (i.e. their capability to infer possibilities from their knowledge). Such a simplified protocol is illustrated in [8] through the interactions between a set of controlled agents and several autonomous uncontrolled ones. These multi-agents plans may be exchanged until a common plan is found [19].

Another case which requires reasoning about the reasoning capabilities of other agents is plan recognition [6]. In both cases, a model of the possible interactions (both positive and negative) between various possible plans is necessary. A taxonomy of such interactions is described in [17]. An architecture to produce plans dealing with time and goal interactions is described in [15].

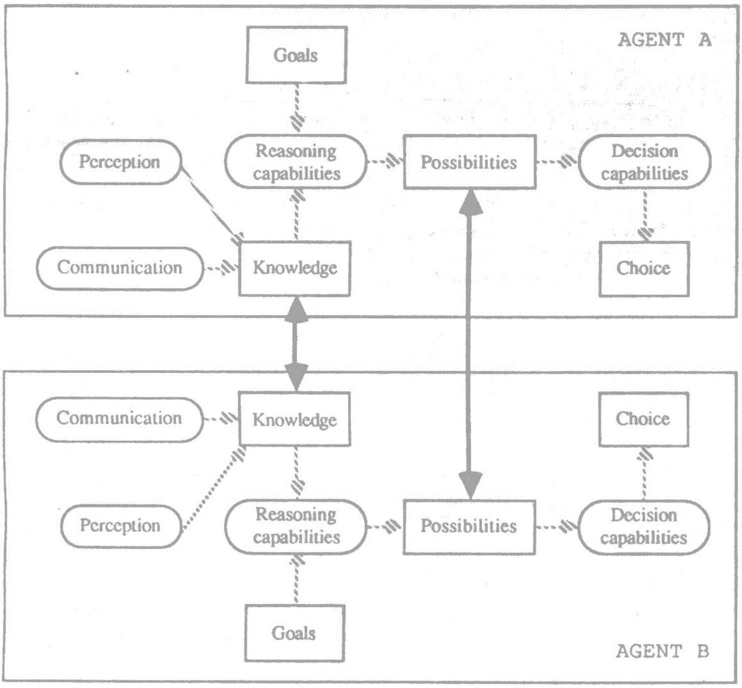

Figure 3: Medium interaction between reasoning capabilities

4.3 Weak interaction between perceiving capabilities

The weakest interaction which can be found among several agents is to use the external world to exchange information. In fact, as we said it before, behaviour can be described as deriving directly or not, from the knowledge an agent has about its environment. Sending information to an agent (communication case) or modifying the environment (world used as a blackboard) are two means by which cooperative behaviour can be achieved among loosely coupled autonomous entities [21]. In fact, no common plans have to be exchanged or even derived explicitly since they do not need to be communicated and no decisions have to be explained because convincing other agents is not part of this paradigm. Another example of this paradigm is in [4] where the environment has been extended to include utterances but in which the agents are described as reactive to changes (situated actions).

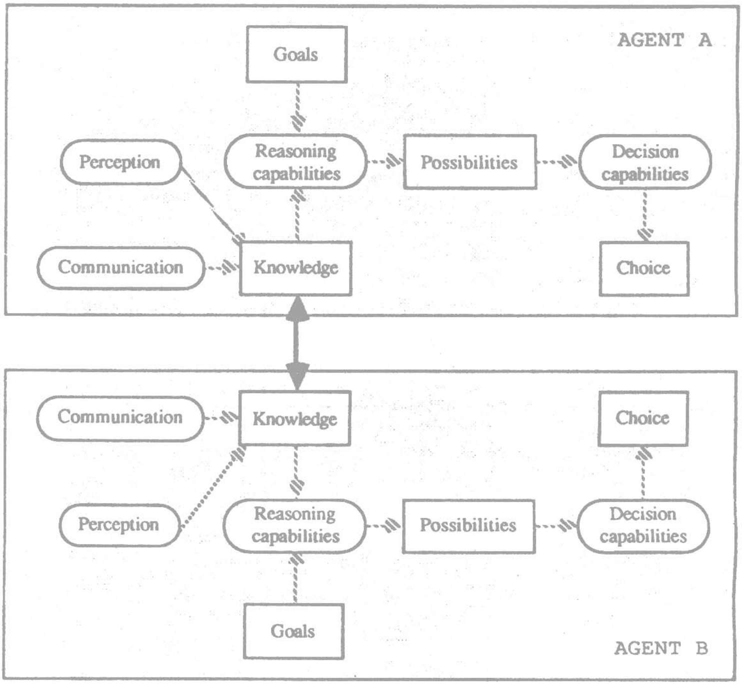

Figure 4: Weak interaction between perceiving capabilities

4.4 Practical considerations for information exchange

All these considerations need strong computational models and a lot of work is focused on providing a framework for designing communicating agents. Two paradigms emerge :

- the *actor* paradigm
- the *blackboard* paradigm

The actor paradigm is based on an object-oriented language where each object has an independent life and communicates with others by sending messages [12]. Various kinds of message passing mechanisms can be imagined as direct addressing, local or global broadcasting [18]. An extension allows typed messages by using abstract descriptions of their structure [11]. The implementation can be either a single machine or distributed.

The blackboard paradigm makes agents communicate by writing on a shared structure called a blackboard. The blackboard can be structured for organizing communication at various levels of abstraction. An agent communicates with another one by writing on the blackboard. Agents are activated when given patterns of information are present on the blackboard [16].

5 Conclusion

By starting from a simple agent description, we have covered a substantial part of possible social behaviours. The consideration of both the locality of a task and the capabilities of an agent to achieve a task, allows us to distinguish four categories of behaviour : cohabitation, cooperation, collaboration and distribution. The consideration of the kinds of information which can be exchanged between several agents seems to allow a fine-grained description of the multi-agent interactions. These results suggest that Decentralized Artificial Intelligence, by emphasizing agent modelling, offers an interesting and systematic framework for studying multi-agent problems.

6 Acknowledgements

Recent developments in DzAI have been and are being performed through an European Action, the Project MAGMA (Modelling an Autonomous aGent in a Multi-Agent world) [5]. This project is part of the COST-13 action of the Commission of the European Communities. The main goal of the program is to study a system architecture, reasoning models and knowledge representation schemes in order to allow problem solving or planning in a multi-agent world. Clearly, this program does not deal with developing a centralized problem solver or planner that will distribute sub-problems or actions to be performed among the agents ; it is concerned with developing the knowledge representation capabilities, the reasoning capabilities, and the decision capabilities which make each agent autonomous in a complex and dynamic multi-agent world. Mobile robots evolving in restricted office environments have been selected as a testbed. We want to thank Pr. Jean-Claude Latombe who inspired this work and the Commission of the European Communities for taking the decision to support it.