


Danny Weyns
H. Van Dyke Parunak
Fabien Michel (Eds.)

Environments for Multi-Agent Systems

First International Workshop, E4MAS 2004

New York, NY, USA, July 2004

Revised Selected Papers



Springer

Tp18-53
E61
2004

Danny Weyns H. Van Dyke Parunak
Fabien Michel (Eds.)

Environments for Multi-Agent Systems

First International Workshop, E4MAS 2004
New York, NY, July 19, 2004
Revised Selected Papers



E200500886



Springer

Series Editors

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

Volume Editors

Danny Weyns
Katholieke Universiteit Leuven
Department of Computer Science, AgentWise, DistriNet
3001 Leuven, Belgium
E-mail: danny.weyns@cs.kuleuven.ac.be

H. Van Dyke Parunak
Altarum Institute
3520 Green Court, Suite 300, Ann Arbor, MI 48105-1579, USA
E-mail: van.parunak@altarum.org

Fabien Michel
LIRMM - Université de Montpellier II
161 rue Ada, 34592 Cedex 5, Montpellier, France
E-mail: fmichel@lirmm.fr

Library of Congress Control Number: 2005920458

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

ISSN 0302-9743
ISBN 3-540-24575-8 Springer 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. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11389903 06/3142 5 4 3 2 1 0

Lecture Notes in Artificial Intelligence

3374

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

Preface

The modern field of multiagent systems has developed from two main lines of earlier research.

Its practitioners generally regard it as a form of artificial intelligence (AI). Some of its earliest work was reported in a series of workshops in the US dating from 1980, revealingly entitled, “Distributed Artificial Intelligence,” and pioneers often quoted a statement attributed to Nils Nilsson that “all AI is distributed.” The locus of classical AI was what happens in the head of a single agent, and much MAS research reflects this heritage with its emphasis on detailed modeling of the mental state and processes of individual agents. From this perspective, intelligence is ultimately the purview of a single mind, though it can be amplified by appropriate interactions with other minds. These interactions are typically mediated by structured protocols of various sorts, modeled on human conversational behavior.

But the modern field of MAS was not born of a single parent. A few researchers have persistently advocated ideas from the field of artificial life (ALife). These scientists were impressed by the complex adaptive behaviors of communities of animals (often extremely simple animals, such as insects or even microorganisms). The computational models on which they drew were often created by biologists who used them not to solve practical engineering problems but to test their hypotheses about the mechanisms used by natural systems. In the artificial life model, intelligence need not reside in a single agent, but emerges at the level of the community from the nonlinear interactions among agents. Because the individual agents are often subcognitive, their interactions cannot be modeled by protocols that presume linguistic competence. The French biologist Grassé observed that these interactions are typically achieved indirectly, through modifications of a shared environment [1].

All interaction among agents of any sort requires an environment. For an AI agent whose interactions with other agents are based on speech act theory, the environment consists of a computer network that can convey messages from one agent’s outbox to another agent’s inbox. For an ALife agent, the environment is whatever the agent’s sensors sense and whatever its effectors try to manipulate.

In most cases, AI agents (and their designers) can take the environment for granted. Error-correcting protocols ensure that messages once sent will arrive in due course. Message latency may lead to synchronization issues among agents, but these issues can be discussed entirely at the level of the agents themselves, without reasoning about the environment. As a result, the environment fades into the background, and becomes invisible.

Not so for ALife agents. Simon observed long ago that the complex behavior of an ant wandering along the ground is best explained not by what goes on inside the ant, but by what happens outside, in the structure of the ground over which the ant moves [2]. When a termite interacts with other termites by depositing

and sensing pheromones, the absorption and evaporation of the pheromone by the environment plays a critical role in the emergent structure of the colony's behavior. There are no error-correcting protocols to ensure that an agent who tries to push a rock from one place to another will in fact be able to realize that objective. From the ALife perspective, the environment is an active participant in agent dynamics, a first-class member of the overall system.

One happy result of the confluence of AI and ALife in MAS is the emergence of hybrid agents that draw on the best of both earlier traditions. This volume, and the workshop of which it is the archival record, is evidence of that hybridization. The agents described in these papers are not artificial ants constructed to test a biologist's theories about insect behavior, but components of systems engineered to fly airplanes, or analyze sensor data, or produce plausible human-like behavior in a video game. Like AI agents, many of them have cognitive, symbolic internal structures. Like ALife agents, all of them interact explicitly and deliberately with the environment through which they coordinate their behaviors.

The notion of the environment in MAS is still young, but the number of papers contributed to this volume suggests the potential of this concept for engineered systems, and their breadth sketches the broad framework of what a mature discipline of environments for multiagent systems might resemble. The entire life cycle of environmental engineering is represented here: conceptual models and languages for the design and specification of environments, simulation environments that admit environments as first-class objects, analysis of the role played by an explicit environment in agent coordination, and examples of full applications that exploit the power of an active environment. The introductory survey pulls these themes together to offer an integrated overview of this emerging discipline.

This volume shows the wide range of exploration typical of a nascent discipline as pioneers discover the best ways to frame problems and approach solutions. It will enable other researchers to take build on this body of initial exploration, and should form the foundation for a fruitful new set of tools and methods for developing multiagent systems.

- [1] Grassé, P.P.: *La Reconstruction du Nid et les Coordinations Inter-individuelles chez Bellicositermes Natalensis et Cubitermes sp. La théorie de la Stigmergie: Essai d'Interprétation du Comportement des Termites Constructeurs.* Insectes Sociaux 6 (1959) 41-84
- [2] Simon, H.A.: *The Sciences of the Artificial.* Cambridge, MA, MIT Press (1969)

Organization

E4MAS 2004 was organized in conjunction with the 3rd International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2004), New York City, July 19, 2004.

Program Co-chairs

Danny Weyns
H. Van Dyke Parunak
Fabien Michel

Katholieke Universiteit Leuven, Belgium
Altarum Institute, Ann Arbor, USA
Université de Montpellier II, France

Program Committee

Eric Bonabeau
Sven Brueckner
Paolo Ciancarini
Yves Demazeau
Marco Dorigo
Alexis Drogoul
Jacques Ferber
Tom Holvoet
Franziska Klügl
Marco Mamei
Jean-Pierre Müller
James Odell
H. Van Dyke Parunak
Karl Tuyls
Paul Valckenaers
Franco Zambonelli

Icosystem, Boston, USA and Paris, France
Altarum Institute, Ann Arbor, USA
University of Bologna, Italy
Laboratoire Leibniz, IMAG, Grenoble, France
Université Libre de Bruxelles, Belgium
Laboratoire d'Informatique de Paris 6, France
Université de Montpellier II, LIRMM, France
DistriNet, K.U.Leuven, Belgium
University of Wurzburg, Germany
University of Modena and Reggio Emilia, Italy
CIRAD-LIRMM, Montpellier, France
James Odell Associates, Ann Arbor, USA
Altarum Institute, Ann Arbor, USA
COMO Lab, V.U.B., Belgium
PMA, K.U.Leuven, Belgium
University of Modena and Reggio Emilia, Italy

Website

<http://www.cs.kuleuven.ac.be/~distrinet/events/e4mas/>

Acknowledgements

Many thanks to the PC members for their critical review work. We also thank Elke Steegmans, Alexander Helleboogh, Kurt Schelfhout, Tom De Wolf, Koen Mertens, Nelis Boucké and Tom Holvoet for their efforts for E4MAS. A special word of thanks to Tom De Wolf for managing the web site.

Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 3403: B. Ganter, R. Godin (Eds.), *Formal Concept Analysis*. XI, 419 pages. 2005.
- Vol. 3398: D.-K. Baik (Ed.), *Systems Modeling and Simulation: Theory and Applications*. XIV, 733 pages. 2005.
- Vol. 3397: T.G. Kim (Ed.), *Artificial Intelligence and Simulation*. XV, 711 pages. 2005.
- Vol. 3374: D. Weyns, H.V.D. Parunak, F. Michel (Eds.), *Environment for Multi-Agent Systems*. X, 279 pages. 2005.
- Vol. 3366: I. Rahwan, P. Moraitis, C. Reed (Eds.), *Argumentation in Multi-Agent Systems*. XII, 263 pages. 2005.
- Vol. 3359: G. Grieser, Y. Tanaka (Eds.), *Intuitive Human Interfaces for Organizing and Accessing Intellectual Assets*. XIV, 257 pages. 2005.
- Vol. 3346: R.H. Bordini, M. Dastani, J. Dix, A.E.F. Seghrouchni (Eds.), *Programming Multi-Agent Systems*. XIV, 249 pages. 2005.
- Vol. 3345: Y. Cai (Ed.), *Ambient Intelligence for Scientific Discovery*. XII, 311 pages. 2005.
- Vol. 3339: G.I. Webb, X. Yu (Eds.), *AI 2004: Advances in Artificial Intelligence*. XXII, 1272 pages. 2004.
- Vol. 3336: D. Karagiannis, U. Reimer (Eds.), *Practical Aspects of Knowledge Management*. X, 523 pages. 2004.
- Vol. 3327: Y. Shi, W. Xu, Z. Chen (Eds.), *Data Mining and Knowledge Management*. XIII, 263 pages. 2005.
- Vol. 3315: C. Lemaître, C.A. Reyes, J.A. González (Eds.), *Advances in Artificial Intelligence – IBERAMIA 2004*. XX, 987 pages. 2004.
- Vol. 3303: J.A. López, E. Benfenati, W. Dubitzky (Eds.), *Knowledge Exploration in Life Science Informatics*. X, 249 pages. 2004.
- Vol. 3275: P. Perner (Ed.), *Advances in Data Mining*. VIII, 173 pages. 2004.
- Vol. 3265: R.E. Frederking, K.B. Taylor (Eds.), *Machine Translation: From Real Users to Research*. XI, 392 pages. 2004.
- Vol. 3264: G. Paliouras, Y. Sakakibara (Eds.), *Grammatical Inference: Algorithms and Applications*. XI, 291 pages. 2004.
- Vol. 3259: J. Dix, J. Leite (Eds.), *Computational Logic in Multi-Agent Systems*. XII, 251 pages. 2004.
- Vol. 3257: E. Motta, N.R. Shadbolt, A. Stutt, N. Gibbins (Eds.), *Engineering Knowledge in the Age of the Semantic Web*. XVII, 517 pages. 2004.
- Vol. 3249: B. Buchberger, J.A. Campbell (Eds.), *Artificial Intelligence and Symbolic Computation*. X, 285 pages. 2004.
- Vol. 3248: K.-Y. Su, J.-i. Tsujii, J.-H. Lee, O.Y. Kwong (Eds.), *Natural Language Processing – IJCNLP 2004*. XVIII, 817 pages. 2005.
- Vol. 3245: E. Suzuki, S. Arikawa (Eds.), *Discovery Science*. XIV, 430 pages. 2004.
- Vol. 3244: S. Ben-David, J. Case, A. Maruoka (Eds.), *Algorithmic Learning Theory*. XIV, 505 pages. 2004.
- Vol. 3238: S. Biundo, T. Frühwirth, G. Palm (Eds.), *KI 2004: Advances in Artificial Intelligence*. XI, 467 pages. 2004.
- Vol. 3230: J.L. Vicedo, P. Martínez-Barco, R. Muñoz, M. Saiz Noeda (Eds.), *Advances in Natural Language Processing*. XII, 488 pages. 2004.
- Vol. 3229: J.J. Alferes, J. Leite (Eds.), *Logics in Artificial Intelligence*. XIV, 744 pages. 2004.
- Vol. 3228: M.G. Hinchey, J.L. Rash, W.F. Truszkowski, C.A. Rouff (Eds.), *Formal Approaches to Agent-Based Systems*. VIII, 290 pages. 2004.
- Vol. 3215: M.G. Negoita, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part III*. LVII, 906 pages. 2004.
- Vol. 3214: M.G. Negoita, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part II*. LVIII, 1302 pages. 2004.
- Vol. 3213: M.G. Negoita, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part I*. LVIII, 1280 pages. 2004.
- Vol. 3209: B. Berendt, A. Hotho, D. Mladenic, M. van Someren, M. Spiliopoulou, G. Stumme (Eds.), *Web Mining: From Web to Semantic Web*. IX, 201 pages. 2004.
- Vol. 3206: P. Sojka, I. Kopeček, K. Pala (Eds.), *Text, Speech and Dialogue*. XIII, 667 pages. 2004.
- Vol. 3202: J.-F. Boulicaut, F. Esposito, F. Giannotti, D. Pedreschi (Eds.), *Knowledge Discovery in Databases: PKDD 2004*. XIX, 560 pages. 2004.
- Vol. 3201: J.-F. Boulicaut, F. Esposito, F. Giannotti, D. Pedreschi (Eds.), *Machine Learning: ECML 2004*. XVIII, 580 pages. 2004.
- Vol. 3194: R. Camacho, R. King, A. Srinivasan (Eds.), *Inductive Logic Programming*. XI, 361 pages. 2004.
- Vol. 3192: C. Bussler, D. Fensel (Eds.), *Artificial Intelligence: Methodology, Systems, and Applications*. XIII, 522 pages. 2004.
- Vol. 3191: M. Klusch, S. Ossowski, V. Kashyap, R. Unland (Eds.), *Cooperative Information Agents VIII*. XI, 303 pages. 2004.
- Vol. 3187: G. Lindemann, J. Denzinger, I.J. Timm, R. Unland (Eds.), *Multiagent System Technologies*. XIII, 341 pages. 2004.

- Vol. 3176: O. Bousquet, U. von Luxburg, G. Rätsch (Eds.), *Advanced Lectures on Machine Learning*. IX, 241 pages. 2004.
- Vol. 3171: A.L.C. Bazzan, S. Labidi (Eds.), *Advances in Artificial Intelligence – SBIA 2004*. XVII, 548 pages. 2004.
- Vol. 3159: U. Visser, *Intelligent Information Integration for the Semantic Web*. XIV, 150 pages. 2004.
- Vol. 3157: C. Zhang, H. W. Guesgen, W.K. Yeap (Eds.), *PRICAI 2004: Trends in Artificial Intelligence*. XX, 1023 pages. 2004.
- Vol. 3155: P. Funk, P.A. González Calero (Eds.), *Advances in Case-Based Reasoning*. XIII, 822 pages. 2004.
- Vol. 3139: F. Iida, R. Pfeifer, L. Steels, Y. Kuniyoshi (Eds.), *Embodied Artificial Intelligence*. IX, 331 pages. 2004.
- Vol. 3131: V. Torra, Y. Narukawa (Eds.), *Modeling Decisions for Artificial Intelligence*. XI, 327 pages. 2004.
- Vol. 3127: K.E. Wolff, H.D. Pfeiffer, H.S. Delugach (Eds.), *Conceptual Structures at Work*. XI, 403 pages. 2004.
- Vol. 3123: A. Belz, R. Evans, P. Piwek (Eds.), *Natural Language Generation*. X, 219 pages. 2004.
- Vol. 3120: J. Shawe-Taylor, Y. Singer (Eds.), *Learning Theory*. X, 648 pages. 2004.
- Vol. 3097: D. Basin, M. Rusinowitch (Eds.), *Automated Reasoning*. XII, 493 pages. 2004.
- Vol. 3071: A. Omicini, P. Petta, J. Pitt (Eds.), *Engineering Societies in the Agents World*. XIII, 409 pages. 2004.
- Vol. 3070: L. Rutkowski, J. Siekmann, R. Tadeusiewicz, L.A. Zadeh (Eds.), *Artificial Intelligence and Soft Computing – ICAISC 2004*. XXV, 1208 pages. 2004.
- Vol. 3068: E. André, L. Dybkjær, W. Minker, P. Heisterkamp (Eds.), *Affective Dialogue Systems*. XII, 324 pages. 2004.
- Vol. 3067: M. Dastani, J. Dix, A. El Fallah-Seghrouchni (Eds.), *Programming Multi-Agent Systems*. X, 221 pages. 2004.
- Vol. 3066: S. Tsumoto, R. Słowiński, J. Komorowski, J.W. Grzymala-Busse (Eds.), *Rough Sets and Current Trends in Computing*. XX, 853 pages. 2004.
- Vol. 3065: A. Lomuscio, D. Nute (Eds.), *Deontic Logic in Computer Science*. X, 275 pages. 2004.
- Vol. 3060: A.Y. Tawfik, S.D. Goodwin (Eds.), *Advances in Artificial Intelligence*. XIII, 582 pages. 2004.
- Vol. 3056: H. Dai, R. Srikant, C. Zhang (Eds.), *Advances in Knowledge Discovery and Data Mining*. XIX, 713 pages. 2004.
- Vol. 3055: H. Christiansen, M.-S. Hacid, T. Andreassen, H.L. Larsen (Eds.), *Flexible Query Answering Systems*. X, 500 pages. 2004.
- Vol. 3048: P. Faratin, D.C. Parkes, J.A. Rodríguez-Aguilar, W.E. Walsh (Eds.), *Agent-Mediated Electronic Commerce V*. XI, 155 pages. 2004.
- Vol. 3040: R. Conejo, M. Urretavizcaya, J.-L. Pérez-de-la-Cruz (Eds.), *Current Topics in Artificial Intelligence*. XIV, 689 pages. 2004.
- Vol. 3035: M.A. Wimmer (Ed.), *Knowledge Management in Electronic Government*. XII, 326 pages. 2004.
- Vol. 3034: J. Favela, E. Menasalvas, E. Chávez (Eds.), *Advances in Web Intelligence*. XIII, 227 pages. 2004.
- Vol. 3030: P. Giorgini, B. Henderson-Sellers, M. Winikoff (Eds.), *Agent-Oriented Information Systems*. XIV, 207 pages. 2004.
- Vol. 3029: B. Orchard, C. Yang, M. Ali (Eds.), *Innovations in Applied Artificial Intelligence*. XXI, 1272 pages. 2004.
- Vol. 3025: G.A. Vouros, T. Panayiotopoulos (Eds.), *Methods and Applications of Artificial Intelligence*. XV, 546 pages. 2004.
- Vol. 3020: D. Polani, B. Browning, A. Bonarini, K. Yoshida (Eds.), *RoboCup 2003: Robot Soccer World Cup VII*. XVI, 767 pages. 2004.
- Vol. 3012: K. Kurumatani, S.-H. Chen, A. Ohuchi (Eds.), *Multi-Agents for Mass User Support*. X, 217 pages. 2004.
- Vol. 3010: K.R. Apt, F. Fages, F. Rossi, P. Szeredi, J. Vánca (Eds.), *Recent Advances in Constraints*. VIII, 285 pages. 2004.
- Vol. 2990: J. Leite, A. Omicini, L. Sterling, P. Torroni (Eds.), *Declarative Agent Languages and Technologies*. XII, 281 pages. 2004.
- Vol. 2980: A. Blackwell, K. Marriott, A. Shimojima (Eds.), *Diagrammatic Representation and Inference*. XV, 448 pages. 2004.
- Vol. 2977: G. Di Marzo Serugendo, A. Karageorgos, O.F. Rana, F. Zambonelli (Eds.), *Engineering Self-Organising Systems*. X, 299 pages. 2004.
- Vol. 2972: R. Monroy, G. Arroyo-Figueroa, L.E. Sucar, H. Sossa (Eds.), *MICAI 2004: Advances in Artificial Intelligence*. XVII, 923 pages. 2004.
- Vol. 2969: M. Nickles, M. Rovatsos, G. Weiss (Eds.), *Agents and Computational Autonomy*. X, 275 pages. 2004.
- Vol. 2961: P. Eklund (Ed.), *Concept Lattices*. IX, 411 pages. 2004.
- 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*. X, 301 pages. 2004.
- Vol. 2930: F. Winkler (Ed.), *Automated Deduction in Geometry*. VII, 231 pages. 2004.
- Vol. 2926: L. van Elst, V. Dignum, A. Abecker (Eds.), *Agent-Mediated Knowledge Management*. XI, 428 pages. 2004.
- Vol. 2923: V. Lifschitz, I. Niemelä (Eds.), *Logic Programming and Nonmonotonic Reasoning*. IX, 365 pages. 2003.
- Vol. 2915: A. Camurri, G. Volpe (Eds.), *Gesture-Based Communication in Human-Computer Interaction*. XIII, 558 pages. 2004.
- Vol. 2913: T.M. Pinkston, V.K. Prasanna (Eds.), *High Performance Computing – HIPC 2003*. XX, 512 pages. 2003.
- Vol. 2903: T.D. Gedeon, L.C.C. Fung (Eds.), *AI 2003: Advances in Artificial Intelligence*. XVI, 1075 pages. 2003.
- Vol. 2902: F.M. Pires, S.P. Abreu (Eds.), *Progress in Artificial Intelligence*. XV, 504 pages. 2003.
- Vol. 2892: F. Dau, *The Logic System of Concept Graphs with Negation*. XI, 213 pages. 2003.

Table of Contents

Survey

Environments for Multiagent Systems State-of-the-Art and Research Challenges

Danny Weyns, H. Van Dyke Parunak, Fabien Michel, Tom Holvoet, Jacques Ferber 1

Conceptual Models

AGRE: Integrating Environments with Organizations

Jacques Ferber, Fabien Michel, José Baez 48

From Reality to Mind: A Cognitive Middle Layer of Environment Concepts for Believable Agents

Paul Hsueh-Min Chang, Kuang-Tai Chen, Yu-Hung Chien, Edward Kao, Von-Wun Soo 57

A Spatially Dependent Communication Model for Ubiquitous Systems

Stefania Bandini, Sara Manzoni, Giuseppe Vizzari 74

Languages for Design and Specification

ELMS: An Environment Description Language for Multi-agent Simulation

Fabio Y. Okuyama, Rafael H. Bordini, Antônio Carlos da Rocha Costa 91

MIC*: A Deployment Environment for Autonomous Agents

Abdelkader Gouaïch, Fabien Michel, Yves Guiraud 109

Simulation and Environments

About the Role of the Environment in Multi-agent Simulations

Franziska Klügl, Manuel Fehler, Rainer Herrler 127

Modelling Environments for Distributed Simulation

Michael Lees, Brian Logan, Rob Minson, Ton Oguara, Georgios Theodoropoulos 150

Mediated Coordination

Supporting Context-Aware Interaction in Dynamic Multi-agent Systems <i>Christine Julien, Gruia-Catalin Roman</i>	168
Environment-Based Coordination Through Coordination Artifacts <i>Alessandro Ricci, Mirko Viroli, Andrea Omicini</i>	190
“Exhibitionists” and “Voyeurs” Do It Better: A Shared Environment for Flexible Coordination with Tacit Messages <i>Luca Tummolini, Cristiano Castelfranchi, Alessandro Ricci, Mirko Viroli, Andrea Omicini</i>	215

Applications

Swarming Distributed Pattern Detection and Classification <i>Sven A. Brueckner, H. Van Dyke Parunak</i>	232
Digital Pheromones for Coordination of Unmanned Vehicles <i>H. Van Dyke Parunak, Sven A. Brueckner, John Sauter</i>	246
Motion Coordination in the Quake 3 Arena Environment: A Field-Based Approach <i>Marco Mamei, Franco Zambonelli</i>	264
Author Index	279

Environments for Multiagent Systems State-of-the-Art and Research Challenges

Danny Weyns¹, H. Van Dyke Parunak², Fabien Michel³,
Tom Holvoet¹, and Jacques Ferber³

¹ AgentWise, DistriNet, K.U.Leuven, B-3001 Leuven, Belgium
{danny.weyns, tom.holvoet}@cs.kuleuven.ac.be

² Altarum Institute, Ann Arbor, MI 48105-1579, USA
van.parunak@altarum.org

³ LIRMM, CNRS, Montpellier, 34392 Montpellier Cedex 5, France
{fmichel, ferber}@lirmm.fr

Abstract. It is generally accepted that the environment is an essential compound of multiagent systems (MASs). Yet the environment is typically assigned limited responsibilities, or even neglected entirely, overlooking a rich potential for the paradigm of MASs.

Opportunities that environments offer, have mostly been researched in the domain of situated MASs. However, the complex principles behind the concepts and responsibilities of the environment and the interplay between agents and environment are not yet fully clarified.

In this paper, we first give an overview of the state-of-the-art on environments in MASs. The survey discusses relevant research tracks on environments that have been explored so far. Each track is illustrated with a number of representative contributions by the research community. Based on this study and the results of our own research, we identify a set of core concerns for environments that can be divided in two classes: concerns related to the structure of the environment, and concerns related to the activity in the environment. To conclude, we list a number of research challenges that, in our opinion, are important for further research on environments for MAS.

1 Introduction

There is a general agreement in the multiagent research community that environments are essential for multiagent systems (MASs). Yet most researchers neglect to integrate the environment as a primary abstraction in models and tools for MASs, or minimize its responsibilities. As a consequence, a rich potential of applications and techniques that can be developed using MASs is overlooked.

Popular frameworks such as Jade [9], Jack [44], Retsina [79] or Zeus [58] reduce the environment to a message transport system or broker infrastructure. Well-known methodologies such as Message [25], Prometheus [66] or Tropos [12] offer support for some basic elements of the environment, however they fail to consider the environment as a first-class entity. Standard literature on MASs

used for education, including [73, 93, 45], only deals very briefly with the topic of environments. Even in the FIPA [34] specifications it is hard to find any functionality for the environment beyond message transport or broker systems. Restricting interaction to inter-agent communication neglects a rich potential of possibilities for the paradigm of MASs.

Researchers working in the domain of situated MASs traditionally integrate the environment as a first-class entity in a MAS. In situated MASs, the environment is an active entity with its own processes that can change its own state, independent of the activity of the embedded agents. Inspired by biological systems, several researchers have shown that the environment can serve as a robust, self-revising, shared memory for agents. This can unburden the individual agents from continuously keeping track of their knowledge about the system. Moreover, it enables the agents to use their environment as an excellent medium for indirect coordination. Gradient fields and evaporating marks in the environment can guide agents in their local context and as such facilitate the coordination in a community of agents in a decentralized fashion. Several practical applications have shown how the environment can contribute to manage complex problems. There are examples in domains such as supply chain systems, network support, peer-to-peer systems, manufacturing control, multiagent simulation etc. Since the exploitation of the environment in MASs results in better manageable solutions, it is a promising paradigm to deal with the increasing complexity and dynamism of future system infrastructure and more advanced problem domains, e.g. ad hoc networks or ubiquitous computing.

Despite the large amount of work in the domain of situated MASs, we are just at the very beginning of understanding the complex principles behind the concepts related to the environment and the interplay between agents and the environment. This paper aims to contribute in three ways. First we give an overview of the state-of-the-art on environments for MASs. Based on this study as well as the results of our own research, we identify a set of core concerns for environments, as a second contribution. Third, we outline a number of research challenges that, in our opinion, are important for the future development of environments for MASs.

2 Organization of the Paper

In Sect. 3, we start with an overview of the state-of-the-art on environments for MASs. Studying MAS literature with a focus on environments is a tough job. During our study, we encountered two types of difficulties: (1) the term *environment* has several different meanings, causing a lot of confusion, (2) the functionalities associated with the environment are often treated implicitly, or integrated in the MAS in an ad-hoc manner.

The confusion on what the environment comprises is mainly caused by mixing up concepts and infrastructure. Sometimes, researchers refer to the environment as the logical entity of a MAS in which the agents and other objects/resources are embedded. Sometimes, the notion of environment is used to refer to the

software infrastructure on which the MAS is executed. Sometimes, environment even refers to the underlying hardware infrastructure on which the MAS runs.

The fact that functionalities of the environment are often treated implicitly, or in an ad-hoc manner, indicates that in general, the MAS research community fails to treat the environment as a *first-class entity*. [36] defines a first-class module as: “a program building block, an independent piece of software which [...] provides an abstraction or information hiding mechanism so that a module’s implementation can be changed without requiring any change to other modules.” Thus, the environment is in general not treated as an independent building block that encapsulates its own clear-cut responsibilities in the MAS, irrespective of the agents.

Starting from this perspective, the overview of the state-of-the-art on environments for MASs we discuss in Sect. 3 is not just a summary of representative papers on the topic of environments for MASs. In fact, the number of research papers that are devoted to the environment is very limited. The overview is rather a reflection on MAS research literature in which we have put the spotlight on models and concepts associated with the environment. The survey is structured as follows:

- 3.1 General models for environments (Russell and Norvig, Ferber, Odell et al., Environments for mobile agents)
- 3.2 Inter-agent facilities
 - Communication infrastructure (Huhns & Stephens, FIPA, Jade, Retsina)
 - Models for indirect interaction
 - Classical blackboard communication
 - Tuple-based interaction models (JavaSpaces, Lime)
 - Stigmergy (Synthetic ecosystems, Network routing)
 - Interaction models related to space in MASs (MMASS)
 - Environment as an organizational layer (AGR)
- 3.3 Agent-Environment interaction
 - Perception of the environment (Robocup Soccer Server, Model for active perception)
 - Dealing with actions in the environment (Synchronous model for action, Action model with regional synchronization)
 - Task-environments (Wooldridge, TAEMS)
- 3.4 Environments in agent-oriented methodologies (Gaia)

For each track we selected a number of relevant contributions from the research community, specified in brackets. It is not a primary goal of the survey to be complete, but rather to give an overview of the wide range of different conceptions associated with the environment for MASs and its various uses.

In Sect. 4, we extract, from the listed research tracks, a set of core concerns for environments in MASs. We have divided the concerns in two classes:

- 4.1 Concerns related to the structure of the environment (Structuring, Resources, Ontology)
- 4.2 Concerns related to the activity in the environment (Communication, Actions, Perception, Environmental processes)

Each concern represents a *logical functionality* for which the environment may have a *natural responsibility*. Our goal is to make the logical functionalities *explicit*, i.e. as concerns of environments as first-class entities. We want to underline that the proposed list of concerns is not intended to be complete. Our aim is to give an initial impetus to explore the rich potential of environments for MASs.

Next in Sect. 5 we outline a number of research challenges that, in our opinion, are important for the further development of environments for MASs. We have divided the list in three categories:

- 5.1 Definition and scope of environments
- 5.2 Agent-environment interrelationship
- 5.3 Engineering environments

Each category discusses a number of applicable research challenges. These challenges may serve as a source of inspiration for future exploration of environments for MASs.

Finally, in Sect. 6 we draw conclusions.

Conventions. In the remainder of the paper, we use the following style conventions:

- Quotations are put in “quotation marks.”
- Specific terms used in literature are marked in *teletype*.
- Terms of concepts we want to emphasize are marked in *italic*.

3 Environments for MASs: A Survey of the State-of-the-Art

In this section we give an overview of a number of important research tracks that, in one way or another, include some notion of environment. We start with discussing a couple of general models for environments that have been proposed in literature. Then we zoom in on various concepts and functionalities related to inter-agent facilities in the environment and agent-environment interaction. We conclude the section by discussing the position of environments in agent-oriented software engineering. Each track is illustrated with a number of relevant contributions from the research community.

3.1 General Models for Environments

We start our study on environments for MASs with a number of representative models for environments that have been proposed in the research community.

Russell and Norvig. In chapter 2 of [73], S. Russell and P. Norvig discuss how an intelligent agent relates to its environment: “An agent is anything that can be viewed as perceiving its environment through sensors and acting upon the environment through effectors.” This generally acknowledged relationship between an agent and its environment is schematically depicted in Fig. 1.




Fig. 1. Agent interaction with the environment [73]

Russell and Norvig discuss a number of key properties of environments that are now adopted by most researchers in the domain:

- Accessible versus inaccessible: indicates whether the agents have access to the complete state of the environment or not.
- Deterministic versus nondeterministic: indicates whether a state change of the environment is uniquely determined by its current state and the actions selected by the agents or not.
- Static versus dynamic: indicates whether the environment can change while an agent deliberates or not.
- Discrete versus continuous: indicates whether the number of percepts and actions are limited or not.

The most complex class of environments are those that are inaccessible, non-deterministic, dynamic and continuous. The first three properties of this list are properties typically occurring in MASSs.

Russell and Norvig also define a “generic environment program”, see Fig. 2. The program periodically gives the agents percepts and receives back their ac-

```

procedure RUN-ENVIRONMENT(state, UPDATE-FN, agents, termination)
  inputs: state, the initial state of the environment
           UPDATE-FN, function to modify the environment
           agents, a set of agents
           termination, a predicate to test when we are done

  repeat
    for each agent in agents do
      PERCEPT[agent] ← GET-PERCEPT(agent, state)
    end
    for each agent in agents do
      ACTION[agent] ← PROGRAM[agent](PERCEPT[agent])
    end
    state ← UPDATE-FN(actions, agents, state)
  until termination(state)
  
```

Fig. 2. A generic environment program [73]

tions. Next, the program updates the state of the environment based on the actions of the agents and of possibly other dynamic processes in the environment that are not considered as agents. This simple program for environments clearly illustrates the basic relationship between agents and their environment.

Ferber. In [28], J. Ferber discusses the modelling of environments for MAS at length. According to Ferber, an environment can either be represented as a single monolithic system, i.e. a centralized environment, or as a set of cells assembled in a network, i.e. a distributed environment. In a centralized environment, all agents have access to the same structure. In a distributed environment, each cell behaves like a centralized environment in miniature. However, a cell of a distributed environment differs in a number of ways from a centralized environment: (1) the state of a cell in a distributed environment depends on the surrounding cells, (2) the perception of agents in a distributed environment typically goes beyond one cell, (3) when agents move from cell to cell, the agent's link with the cells has to be managed and (4) the propagation of signals over the network of cells has to be managed. Orthogonal to the difference between a centralized or a distributed representation of environment, Ferber distinguishes between "generalized" and "specialized" models for environments. A generalized model is independent of the kind of actions that can be performed by agents. A specialized model is characterized by a well-defined set of actions. Ferber further distinguishes between purely communicative MASs (in which agents can only communicate by message transfer), purely situated MASs (in which agents can only act in the environment) and the combination of communicating and situated MASs.

Central to Ferber's model of an environment is the way actions are modelled. The action model of Ferber distinguishes between *influences* and *reactions* to influences. Influences come from inside the agents and are attempts to modify the course of events in the world. Reactions, which result in state changes, are produced by the environment by combining influences of all agents, given the local state of the environment and the laws of the world. This clear distinction between the products of the agents' behavior and the reaction of the environment provides a way to handle simultaneous activity in the MAS.

Ferber uses the BRIC formalism (Block-like Representation of Interactive Components) to model a MAS as a set of interconnected components that can exchange messages via the links. BRIC components encapsulate their own behavior and can be composed hierarchically. Fig. 3 depicts a model for a combined communicating and situated MAS in BRIC notation. In the BRIC model depicted in Fig. 3, the activity cycle of the MAS starts when the environment sends "perceptions" to the agents. As soon as the Synchronizer sends "synchronization of perceptions" signals to the agents, the agents are triggered to interpret the perceptions. Then, each agent produces an influence in the environment and possibly transmits a message to another agent. Next, the agent informs the Synchronizer it has finished its action by sending an "synchronization of actions" message. When all agents have sent their "synchronization of actions" messages, the Synchronizer sends a "synchronization of reactions" message to the Environ-