


Rogier M. van Eijk
Marc-Philippe Huget
Frank Dignum (Eds.)

Agent Communication

International Workshop on Agent Communication, AC 2004
New York, NY, USA, July 2004
Revised Selected and Invited Papers



TP18-53
A265.4
2004

Rogier M. van Eijk Marc-Philippe Huget
Frank Dignum (Eds.)

Agent Communication

International Workshop on Agent Communication, AC 2004
New York, NY, USA, July 19, 2004
Revised Selected and Invited Papers



E200500905



Springer

Series Editors

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

Volume Editors

Rogier M. van Eijk
Frank Dignum
Utrecht University
Institute of Information and Computing Sciences
3508 TB Utrecht, The Netherlands
E-mail: {rogier, dignum}@cs.uu.nl

Marc-Philippe Huget
ESIA-LISTIC
B.P. 806, 74016 Annecy, France
E-mail: Marc-Philippe.Huget@univ-savoie.fr

Library of Congress Control Number: 2005920596

CR Subject Classification (1998): I.2.11, I.2, C.2.4, C.2, D.2, F.3

ISSN 0302-9743

ISBN 3-540-25015-8 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: 11394303 06/3142 5 4 3 2 1 0

Lecture Notes in Artificial Intelligence

3396

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

Preface

In this book, we present a collection of papers around the topic of agent communication. The communication between agents has been one of the major topics of research in multiagent systems. The current work can therefore build on a number of previous Workshops of which the proceedings have been published in earlier volumes in this series. The basis of this collection is formed by the accepted submissions of the Workshop on Agent Communication held in conjunction with the AAMAS Conference in July 2004 in New York. The workshop received 26 submissions of which 14 were selected for publication in this volume. Besides the high-quality workshop papers we noticed that many papers on agent communication found their way to the main conference. We decided therefore to invite a number of authors to revise and extend their papers from this conference and to combine them with the workshop papers. We believe that the current collection comprises a very good and quite complete overview of the state of the art in this area of research and gives a good indication of the topics that are of major interest at the moment.

The papers can roughly be divided over the following five themes:

- social commitments
- multiparty communication
- content languages
- dialogues and conversations
- speech acts

Although these themes are of course not mutually exclusive they indicate some main directions of research. We therefore have arranged the papers in the book according to the topics indicated above.

The first three papers focus on the role of social commitments in agent communication. In the first paper, Nicoletta Fornara, Fransesco Viganò and Marco Colombetti explore the role of social commitments in defining the semantics of agent communication in the context of artificial institutions. *Roberto Flores, Philippe Pasquier and Brahim Chaib-draa* formalize the dynamics of social commitments, where they stress the role of commitment messages as coordination devices to advance the state of joint activities. In the subsequent contribution, *Ashok Mallya and Munindar Singh* use social commitments as a semantic underpinning of a formal framework to reason about the composition of interaction protocols.

The next two contributions involve communication between more than two agents. *Gery Gutnik and Gal Kaminka* address the representation of interaction protocols by means of Petri nets. In particular, the authors focus on protocols for overhearing in which more than two agents are involved. The theme of multiparty communication is further elaborated in the contribution of *Marc-Philippe Huget and Yves Demazeau*, where a communication server for multiparty dialogue is described.

The following two contributions focus on the role of vocabularies, ontologies and content languages in agent communication. *Jurriaan van Diggelen, Robbert-Jan Beun, Frank Dignum, Rogier van Eijk and John-Jules Meyer* study the characteristics and properties of communication vocabularies in multiagent systems with heterogeneous ontologies. *Mario Verdicchio and Marco Colombetti* deal with another aspect of content languages: the formal expression of temporal conditions.

The first paper of the section on dialogues and conversations is by *Jarred McGinnis and David Robertson* who define a general language for the expression of dynamic and flexible dialogue protocols. The flexibility of protocols is further elaborated in the contribution of *Lalana Kagal and Tim Finin* where conversation specifications and conversation policies are defined in terms of permissions and obligations. The authors introduce techniques to resolve conflicts within the specifications and policies and provide an engine that allows agents to reason about their conversations. In the next paper, *Joris Hulstijn, Mehdi Dastani and Frank Dignum* study the coherence of agent conversations. In particular, they show how constraints on the context of messages can be used to establish coherent dialogues. The importance of the context is also stressed in the contribution of *Matthias Nickles, Michael Rovatsos and Gerhard Weiss* who study the effects of social interaction structures on the semantics of messages. *Mirko Viroli and Alessandro Ricci* look at communication from the perspective of coordination. In their approach, agents coordinate their activities via artifacts that specify the successive actions of the interaction protocol.

The last four contributions of the volume are centered around the semantics of speech acts. *Karim Bouzouba, Jamal Bentahar and Bernard Moulin* develop a computation model to study the semantics of speech acts in dialogues between agents and humans. In the subsequent contribution, *Peter McBurney and Simon Parsons* propose a set of speech acts together with an interaction protocol for argumentation for which they provide an operational semantics. In the contribution of *Marcus Huber, Sanjeev Kumar and David McGee*, a repertory of speech acts is provided, where the semantics of the acts is defined in terms of joint intentions. Finally, *Shakil Khan and Yves Lesperance* study the semantics of speech acts in terms of the agents' knowledge, intentions and commitments and show how this can be integrated into a planning framework.

To close, we would like to take this opportunity to thank the members of the Program Committee, the external reviewers and the authors of submitted papers for enabling us to edit this exciting volume on the *Developments in Agent Communication*.

Utrecht, November 2004

Rogier van Eijk
Marc-Philippe Huget
Frank Dignum

Workshop Organization

Organizing Committee

Rogier van Eijk	Utrecht University, Utrecht, The Netherlands
Marc-Philippe Huget	Laboratoire LEIBNIZ, Institut IMAG, France
Frank Dignum	Utrecht University, Utrecht, The Netherlands

Program Committee

Leila Amgoud	IRIT (France)
Brahim Chaib-draa	Laval University (Canada)
Phil Cohen	Oregon Health and Science University (USA)
Marco Colombetti	Politecnico di Milano (Italy)
Mehdi Dastani	Utrecht University (The Netherlands)
Amal El Fallah-Seghrouchni	University of Paris 6 (France)
Frank Guerin	University of Aberdeen (UK)
Mark d'Inverno	Westminster University (UK)
Andrew Jones	King's College, London (UK)
Yannis Labrou	Fujitsu Laboratories (USA)
Nicolas Maudet	University of Paris 9 (France)
Peter McBurney	University of Liverpool (UK)
Simon Parsons	Brooklyn College, City University of NY (USA)
Shamima Paurobally	University of Southampton (UK)
Nico Roos	Maastricht University (The Netherlands)
Munindar Singh	North Carolina State University (USA)
Gerhard Weiss	Technical University Munich (Germany)
Michael Wooldridge	University of Liverpool (UK)

External Reviewers

Sanjeev Kumar
Jan Broersen
Roberto Flores
Jamal Bentahar
Philippe Pasquier
Michael Rovatsos
Matthias Nickles

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. 3396: R.M. van Eijk, M.-P. Huget, F. Dignum (Eds.), *Agent Communication*. X, 261 pages. 2005.
- Vol. 3374: D. Weyns, H.V.D. Parunak, F. Michel (Eds.), *Environments 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. 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. Kopecek, 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

Section I: Social Commitments

Agent Communication and Institutional Reality <i>Nicoletta Fornara, Francesco Viganò, Marco Colombetti</i>	1
Conversational Semantics with Social Commitments <i>Roberto A. Flores, Philippe Pasquier, Brahim Chaib-draa</i>	18
A Semantic Approach for Designing Commitment Protocols <i>Ashok U. Mallya, Munindar P. Singh</i>	33

Section II: Multi-party Communication

A Scalable Petri Net Representation of Interaction Protocols for Overhearing <i>Gery Gutnik, Gal Kaminka</i>	50
First Steps Towards Multi-party Communication <i>Marc-Philippe Huget, Yves Demazeau</i>	65

Section III: Content Languages

Optimal Communication Vocabularies and Heterogeneous Ontologies <i>Jurriaan van Diggelen, Robbert Jan Beun, Frank Dignum, Rogier M. van Eijk, John-Jules Meyer</i>	76
Dealing with Time in Content Language Expressions <i>Mario Verdicchio, Marco Colombetti</i>	91

Section IV: Dialogues and Conversations

Realizing Agent Dialogues with Distributed Protocols <i>Jarred McGinnis, David Robertson</i>	106
Modeling Communicative Behavior Using Permissions and Obligations <i>Lalana Kagal, Tim Finin</i>	120
Coherence Constraints for Agent Interaction <i>Joris Hulstijn, Frank Dignum, Mehdi Dastani</i>	134

Formulating Agent Communication Semantics and Pragmatics as
Behavioral Expectations
Matthias Nickles, Michael Rovatsos, Gerhard Weiss 153

Agent Interaction Semantics by Timed Operating Instructions
Mirko Viroli, Alessandro Ricci 173

Section V: Speech Acts

Dialogization and Implicit Information in an Agent Communicational
Model
Karim Bouzouba, Jamal Bentahar, Bernard Moulin 193

Locutions for Argumentation in Agent Interaction Protocols
Peter McBurney, Simon Parsons 209

Toward a Suite of Performatives Based Upon Joint Intention Theory
Marcus J. Huber, Sanjeev Kumar, David McGee 226

A Model of Rational Agency for Communicating Agents
Shakil M. Khan, Yves Lespérance 242

Author Index 261

Agent Communication and Institutional Reality*

Nicoletta Fornara¹, Francesco Viganò¹, and Marco Colombetti^{1,2}

¹ Università della Svizzera italiana, via G. Buffi 13, 6900 Lugano, Switzerland
{nicoletta.fornara, francesco.vigano, marco.colombetti}@lu.unisi.ch

² Politecnico di Milano, piazza Leonardo Da Vinci 32, Milano, Italy
marco.colombetti@polimi.it

Abstract. In this paper we propose to regard an Agent Communication Language (ACL) as a set of conventions to act on a fragment of institutional reality, defined in the context of an artificial institution. Within such an approach, we first reformulate a previously proposed commitment-based semantics for ACLs. In particular we show that all commonly used types of communicative acts can be defined in terms of a single basic type, namely *declarations*, within an artificial institution that we call Basic Institution. We then go on defining special institutions, that augment the Basic Institution by adding ontological and normative elements. Finally, as an example of a special institution we give a partial definition of the institution of English Auctions.

1 Introduction

In the last few years the concept of *social commitment* has been largely used by a growing number of researchers to define the semantics of Agent Communication Languages (ACLs). After the first studies carried out by Singh and by Colombetti [28, 5], further investigations have been carried out from an operational point of view [12, 19], following a logical approach [30], and in the field of argumentation studies [1, 3]. The main advantages of this approach are that commitments are objective and independent of an agent's internal structure, and that it is possible to verify whether an agent is behaving according to the given semantics.

Social commitments are used to represent the evolution of social relationships among agents during interactions. Communicative acts are then viewed as actions carried out to modify such relationships by creating, updating or cancelling commitments according to a predefined set of shared rules [30, 13]. More precisely, communicative acts are regarded as a sort of *institutional actions*, that is, as actions performed within an *institution* to modify a fragment of social reality [25]. Defining the semantics of an ACL has therefore two sides: one side is the definition of the institutional effects brought about by the performance of communicative acts; the other side is the definition of the social context in which

* Supported by Swiss National Science Foundation project 200021-100260, "An Open Interaction Framework for Communicative Agents".

agents can carry out institutional actions, and that we call an (*artificial*) *institution*. Indeed, our main tenet is that without the definition of an appropriate institution it is impossible to specify the semantics of an ACL.

This paper is structured as follows. In Section 2 we introduce the fundamental concepts on which we base our treatment of agent communication, namely the concepts of an *institutional action*, of a *convention*, and of a "counts as" relationship between an *instrumental action* and the corresponding communicative act. In Section 3 we define the institutional actions that can be performed on commitments. In Section 4 we describe the Basic Institution (i.e., the institution that regulates the management of commitments) and introduce the concept of a special institution. In Section 5 we give a partial description of a specific case of a special institution, that is, the institution of English Auctions. In Section 6 we briefly remark on related work present in this volume. Finally in Section 7 we draw some conclusions and delineate some directions for future work.

2 Fundamental Concepts

We view a multiagent system (MAS) as a technological extension of human society, by which individual persons and human organizations can delegate the execution of institutional actions to the artificial system. Examples of such actions are establishing appointments, signing contracts, and carrying out commercial transactions. For this reason there are strong connections between some aspects of a MAS and some aspects of human society, and therefore the concepts used to model a MAS interaction framework have to reflect some crucial characteristics of their human counterpart.

The context within which artificial agents operate can be modelled as consisting of a set of *entities* that can have *natural* or *institutional* attributes, that is, attributes that exist only thanks to the common agreement of the interacting agents (or more precisely of their users). For example, the color of a book is a natural attribute, while the book's price and its owner are institutional attributes. Natural attributes are assumed to reflect the physical properties of the corresponding entities of the real world, and typically cannot be changed by artificial agents (unless the agent controls a physical robot). On the contrary, institutional attributes can be affected by *institutional actions* performed by purely software agents.

2.1 Institutional Actions

Institutional actions are particular types of actions [7] that are crucial for the formalization of communicative interactions taking place in open interaction frameworks. The effect of institutional actions is to change institutional attributes, that exist only thanks to common agreement. Therefore, agents cannot perform such actions by exploiting causal links occurring in the natural world, as it would be done to open a door or to remove a physical object. Rather, as we shall see, institutional actions are performed on the basis of a shared set of conventions.

Because of their intrinsic social nature, a crucial condition for the actual performance of institutional actions is that they must be *public*, that is, made known to the relevant agents by means of some action that can be directly executed by an artificial agent. It is therefore natural to assume that all institutional actions are performed by sending suitable messages to the relevant agents. An example of institutional action, that will be discussed in Section 5, is the act of opening an auction; as we shall see, an agent (the auctioneer) can perform such an action by sending a suitable message to the relevant group of agents (the participants). However, the act of sending the message is merely instrumental, and should not be confused with the institutional action of opening the auction.

We define institutional actions by specifying their *preconditions* and *postconditions*, therefore abstracting from the way in which such actions are concretely carried out. More precisely, an institutional action is characterized by:

- an *action name* followed by a possibly empty list of *parameters*;
- a possibly empty set of (*ontological*) *preconditions*, that specify the values that certain institutional attributes must have for the action to be meaningful (for example, opening an auction is meaningful only if the auction is not already open);
- a nonempty set of *postconditions*, that specify the values of certain institutional attributes after a successful performance of the action.

2.2 Instrumental Actions

As we have already remarked, an institutional action is performed by executing an instrumental action, conventionally associated to the institutional action. In the human world such instrumental actions vary from certain bodily movements (raising one's arm to vote), to the use of specific physical tools (waving a white flag to surrender), to the use of language (saying "the auction is open" to open an auction). In a system of artificial agents, it is natural to assume that all institutional actions are performed by means of a single type of instrumental actions, namely exchanging a message.

For the purposes of the current treatment, a message consists of: a *message type*, a *sender*, one or more *receivers*, and a *content*. The action of exchanging a message will be represented with the following notation:

`exchMsg(message_type, sender, receiver(s), content)`

Note that here *sender* and *receiver* are just fields of a message. That such fields correctly represent the agent that actually sends the message and the agents to which the message is delivered has to be guaranteed by the underlying message transport system.

2.3 The "Counts as" Relation

Following Searle [25], the construction of social reality in the human world is possible thanks to *constitutive rules* of the form *X counts as Y in C*; in the

particular case where X and Y are actions, the performance of an action of type X in context C can count as performing an action of type Y . Similarly, in an artificial system, thanks to shared *conventions*, the action of exchanging a particular message can "count as" the execution of some institutional action, if certain *contextual conditions* are satisfied.

According to Searle's Speech Act Theory [24], *declarations* are the particular category of communicative acts whose point is to bring about a change in the institutional reality in virtue of their successful performance. By definition the content of a declaration describes precisely the institutional changes that it brings about. Therefore, we take messages of type *declare* as the fundamental means to perform institutional actions. The convention that binds the exchange of a *declare* message to the performance of the institutional action (*iaction*), described in its content, can be described as follows:

$$\text{exchMsg}(\text{declare}, \text{sender}, \text{receiver}, \text{iaction}(\text{parameters})) =_{\text{conv}} \text{iaction}(\text{parameters})$$

By itself, however, a convention is not sufficient to guarantee the successful performance of an institutional action by the exchange of a *declare* message: indeed, some additional conditions about the agent that sends the message and about the state of system must be satisfied. In general, an agent must be *authorized* to perform an institutional action; for example, only the auctioneer can open an auction by sending a suitable message to the participants. Further *contextual conditions* about the state of the system, expressed by suitable Boolean expressions, may be required; for example, it may be established that an auction is validly opened only if there are at least two participants.

Assuming that every agent in the interaction system has an identifier (*agent_id*), authorizations will be represented with the following notation:

$$\text{Auth}(\text{agent_id}, \text{iaction}(\text{parameters}), \text{contextual_conditions})$$

Our notion of authorization should not be confused with the notion of *permission*. The distinction we make between these two concepts is similar to the one between institutionalized power and permission proposed by Jones and Sergot in [16]. While authorizations are necessary conditions for the performance of institutional actions, permissions (like obligations) are brought about by *norms* (see Section 4.2), that is, by rules that affect the normative positions of the agents in the system. The crucial difference between authorizations and permissions is highlighted in the cases when they are not granted. If an agent is not authorized to perform an institutional action, a performance of the corresponding instrumental action does not count as a performance of the institutional action (the institutional action is thus not executed). On the contrary, if an authorized agent performs an institutional action without permission, the institutional action is successfully performed, but the agent violates a norm and it may be sanctioned for its behavior.

In the specification of an interaction system it is useful to express authorizations in term of the *roles* filled by agents, in order to abstract from the concrete

agents that will be actually involved in an interaction. For example, the authorization to open and close an auction is granted to the agent that fills the role of the auctioneer, independently of its individual identity.

The concept of a role is very broad: for example, it is possible to regard social commitments as institutional entities that define two roles: the *debtor* of the commitment and its *creditor*. This fact appears to be general; that is, roles are defined relative to an institutional entity. We can then abstractly define the authorization to perform a specific institutional action (with given parameters) associating it to a role defined in the context of a specific institutional entity (*identity*):

$$\text{Auth}(\text{identity.role}, \text{iaction}(\text{parameters}), \text{contextual_conditions})$$

In a concrete interaction, the authorizations associated to roles need to be transformed into authorizations of an actual agent in the system. Such transformation can be obtained searching among all the institutional entities present in the system the ones that match the description given through the parameters of the institutional action, and then creating a concrete authorization for each agent having the role indicated in the abstract authorization.

3 A Commitment-Based Agent Communication Language

The semantics of ACLs that we have proposed in [12, 13] is based on the assumption that the performance of a communicative act in a multiagent system has the effect of changing the social relationship between the sender and the receiver, and that this change can be represented by means of an institutional entity, that is, *social commitment*. To specify the meaning of various types of communicative acts in terms of effects on commitments, it is necessary to define an ontology of commitment and the institutional actions necessary to operate on commitments.

3.1 The Ontology of Commitment

We regard a commitment as an entity with the following attributes: a *debtor*; a *creditor*; a *content*; a *state*, used to keep track of the temporal evolution of the commitment. Commitments will be represented with the following notation:

$$\text{Comm}(\text{state}, \text{debtor}, \text{creditor}, \text{content})$$

The content of a commitment can be represented by means of a *temporal proposition* (for a detailed treatment of temporal propositions see [13, 6]), that is, a proposition about a state of affairs or about the performance of an action, referred to a specific interval of time. At every time instant, a temporal proposition has a truth value, that can be *undefined*, *true*, or *false*.

The state of a commitment undergoes a life cycle, described by the state diagram of Figure 1, and can change as an effect of the execution of institutional

actions (solid lines) or of environmental events (dotted lines). Relevant events are due to the change of the truth-value of the commitment's content.




Fig. 1. The life-cycle of commitments

The creditor of a commitment can be a single agent or a group of agents. It is important to remark that a commitment taken with a group of agents need not be equivalent to a conjunction of commitments taken with every member of the group. This point has been thoroughly analyzed in the literature [4, 8] but is behind the scope of this paper.

Institutional Actions on Commitment. The institutional actions that operate on commitments are defined below; preconditions and effects are described using Object Constraint Language (OCL) [23].

- *name* : `makeCommitment(debtor, creditor, content)`
pre : $\text{not Comm.allInstances} \rightarrow \text{exists}(c | c.\text{debtor} = \text{debtor} \text{ and } c.\text{creditor} = \text{creditor} \text{ and } c.\text{content} = \text{content})$
post : $\text{Comm.allInstances} \rightarrow \text{exists}(c | c.\text{state} = \text{unset} \text{ and } c.\text{debtor} = \text{debtor} \text{ and } c.\text{creditor} = \text{creditor} \text{ and } c.\text{content} = \text{content})$
- *name* : `setCancel(debtor, creditor, content)`
pre : $\text{Comm.allInstances} \rightarrow \text{exists}(c | (c.\text{state} = \text{unset} \text{ or } c.\text{state} = \text{pending}) \text{ and } c.\text{debtor} = \text{debtor} \text{ and } c.\text{creditor} = \text{creditor} \text{ and } c.\text{content} = \text{content})$
post : $\text{Comm.allInstances} \rightarrow \text{exists}(c | c.\text{state} = \text{cancel} \text{ and } c.\text{debtor} = \text{debtor} \text{ and } c.\text{creditor} = \text{creditor} \text{ and } c.\text{content} = \text{content})$
- *name* : `setPending(debtor, creditor, content)`
pre : $\text{Comm.allInstances} \rightarrow \text{exists}(c | c.\text{state} = \text{unset} \text{ and } c.\text{debtor} = \text{debtor} \text{ and } c.\text{creditor} = \text{creditor} \text{ and } c.\text{content} = \text{content})$
post : $\text{Comm.allInstances} \rightarrow \text{exists}(c | c.\text{state} = \text{pending} \text{ and } c.\text{debtor} = \text{debtor} \text{ and } c.\text{creditor} = \text{creditor} \text{ and } c.\text{content} = \text{content})$

It is often useful to define *institutional macro-actions*, that is, actions whose execution coincides with the sequential execution of a list of existing institutional actions, conceived of as a single transaction. For example:

name: `makePendingComm(debtor, creditor, content) =def makeCommitment(debtor, creditor, content), setPending(debtor, creditor, content)`