

Frank Dignum
Rogier M. van Eijk
Roberto Flores (Eds.)

LNAI 3859

Agent Communication II

International Workshops
on Agent Communication, AC 2005 and AC 2006
Utrecht, Netherlands, July 2005
and Hakodate, Japan, May 2006, Selected and Revised Papers



Springer

TP1853

A265.4

2005-2006

Frank Dignum Rogier M. van Eijk
Roberto Flores (Eds.)

Agent Communication II

International Workshops
on Agent Communication, AC 2005 and AC 2006
Utrecht, Netherlands, July 25, 2005
and Hakodate, Japan, May 9, 2006
Selected and Revised Papers



Springer



E2007001414

Series Editors

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

Volume Editors

Frank Dignum
Department of Information and Computing Sciences
Universiteit Utrecht
3508 TB Utrecht, The Netherlands
E-mail: dignum@cs.uu.nl

Rogier M. van Eijk
Department of Information and Computing Sciences
Universiteit Utrecht
3508 Utrecht, The Netherlands
E-mail: rogier@cs.uu.nl

Roberto Flores
Department of Physics
Computer Science and Engineering
Christopher Newport University
Newport News VA 23606, USA
E-mail: flores@pcs.cnu.edu

Library of Congress Control Number: 2006937885

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

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743
ISBN-10 3-540-68142-6 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-68142-7 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

springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11947745 06/3142 5 4 3 2 1 0

Lecture Notes in Artificial Intelligence 3859

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

Preface

Although everyone recognizes communication as a central concept in multi-agents, many no longer see agent communication as a research topic. Unfortunately there seems to be a tendency to regard communication as a kind of information exchange that can easily be covered using the standard FIPA ACL. However, the papers in this volume show that research in agent communication is far from finished. If we want to develop the full potential of multi-agent systems, agent communication should also develop to a level beyond parameter or value passing as is done in OO approaches!

In this book we present the latest collection of papers around the topic of agent communication. The collection comprises of the best papers from the agent communication workshops of 2005 and 2006, enriched with a few revised agent communication papers from the AAMAS conference. Due to some unfortunate circumstances the proceedings of the 2005 workshop were delayed, but it gave us the opportunity to join the best papers of the 2006 edition to this volume. Together these papers give a very good overview of the state of the art in this area of research and give a good indication of the topics that are of major interest at the moment.

The papers are divided into the following four topics:

- Semantics of Agent Communication
- Commitments in Agent Communication
- Protocols and Strategies
- Reliability and Overhearing

The research on the semantics of agent communication has shifted from concepts based on mental attitudes of the agents towards concepts based on social attitudes. However, FIPA ACL, the de facto agent communication standard language, still has an official semantics based on mental attitudes. The first paper in this volume by *V. Louis and T. Martinez* describes a first attempt to actually develop a tool to support verification of compliance to the FIPA ACL semantics, therefore giving an operational semantics to FIPA ACL. The paper of *U. Endriss* shows how conversation protocols can be described in terms of temporal logic. This allows model checking of the protocols and verifying their correctness with respect to a number of properties. The third paper on the semantics of agent communication by *G. Boello, R. Damiano, J. Hulstijn and L. van der Torre* is one of the first that tries to combine private and social mental attitudes to determine the semantics of the communication. Another effort to combine social commitments with the other attitudes of an agent can be found in the paper by *P. Pasquier and B. Chaib-draa*, which is placed under the second topic of social commitments, but could also have been placed under this heading. The last paper in this section of the volume by *S. Khan and Y. Lesperance* describes a possible semantics for conditional commitments. This is an important step

because since commitments are themselves used as semantics of the communication it is important to have some grasp of the characteristics of this concept as well!

The second section of the volume is completely devoted to papers on social commitments. The first paper in this section by *M. Verdicchio and M. Colombetti* discusses how the concept of commitment can be used to build up a library of speech acts to be used in agent communication. A similar topic is discussed in the paper of *R. Kremer and R. Flores*. They add the idea of organizing the speech acts according to a subsumption hierarchy to facilitate their processing. The paper of *G. Muller and L. Vercouter* discusses the use of social commitments that follow from speech acts as constraints on subsequent communication, therefore truly utilizing the intuitions that come with the “commitment” concept. Once commitments are seen as constraints on the behavior of the agents, one should also consider what should be done if the agents do not comply to these constraints. This aspect is discussed in the paper of *J. Heard and R. Kremer*. The last paper in the section on commitments by *A. Mallya and M. Singh* discusses a more advanced aspect of commitments. It looks at preferences with respect to different commitment protocols which can be used to solve potential conflicts between different commitments.

The third section of this volume is devoted to protocols and strategies for communication. It starts with a paper by *L. Amgoud and S. Kaci* on strategies for agents to be used during negotiation. The idea is to make these strategies less restrictive such that good compromises are not discarded too quickly. The second paper in this section by *J. van Diggelen, E. de Jong and M. Wiering* also discusses strategies for communication, but in the domain of ontology alignment. What is the influence of the strategies that agents use to adopt concepts on the convergence to the use of one or more concepts? An important question in order to see what will happen if thousands of agents with different ontologies on the Web start communicating. The next paper by *J. van Diggelen, R.-J. Beun, F. Dignum, R. van Eijk and J.-J. Meyer* actually continues this discussion. It describes how agents can align their ontologies on the fly using some extensions of normal communication protocols. Instead of learning complete ontologies the agents just learn enough about the concepts to be able to use them properly.

The paper of *Pinar Yolum* discusses the important topic of designing protocols and shows some tools that can support this process. Whereas most approaches see protocols as a kind of finite state machine or Petri net, *F. Fischer, M. Rovatsos, and G. Weiss* see protocols as patterns that can be adapted. The semantics follows from the use rather than the use from the semantics. It is thus a perfect example of the use of a bottom approach for creating semantics to communication.

The last two papers in this section discuss communication in relation to the beliefs of the agents. The first paper, by *H.-J. Lebbink, C. Witteman and J.-J. Meyer*, discusses conversations about changing one’s beliefs. When can an agent conclude that it should retract some beliefs based on information it hears from another agent? Preferably this only happens in a way that keeps the beliefs of

the agents “consistent” as far as possible. The paper by *I. Letia and R. Vartic* discusses the different consequences of basing communication on firm beliefs and on defeasible beliefs. Making this distinction allows for more subtle conversations that also seem to resemble human conversations more closely.

The last section of this volume contains papers that deal with multi-party communication. The paper of *S. Cranefield* discusses group communication in which a reliable shared perception of the order of the messages can be very important. This leads to a design of a type of synchronous group communication. The second paper by *N. Dragoni, M. Gaspari and D. Guidi* discusses the very important issue of communication breakdowns. Especially communication over the Internet should be fault-tolerant in order to work for large applications. They discuss a fault-tolerant ACL and illustrate its use on the Web. The final paper of this volume also discusses reliability of communication. In the approach of *G. Gutnik and G. Kaminka* this is achieved by selective overhearing of communication by other agents. Because trying to monitor all communication is prohibitively expensive they propose a hierarchical organization of the agents that can perform a selective overhearing.

We want to conclude this preface by extending our thanks to the members of the Program Committee of the ACL workshops that were willing to review the papers in a very short time span, to the external reviewers that probably even had less time to review their papers and also of course to the authors that were willing to submit their papers to our workshops and the authors that revised their papers for this book.

September 2006

Frank Dignum (Utrecht, Netherlands)
 Rogier van Eijk (Utrecht, Netherlands)
 Roberto Flores (Newport News, USA)

Workshop Organization

Organizing Committee

Rogier van Eijk	Universiteit Utrecht, Utrecht, Netherlands
Roberto Flores	Christopher Newport University, USA
Marc-Philippe Huget	University of Savoie, Annecy, France
Frank Dignum	Universiteit Utrecht, Netherlands, (2006)

Program Committee

L. Amgoud	IRIT, France
J. Bentahar	Laval University, Canada
B. Chaib-draa	Laval University, Canada
P. Cohen	Oregon Health and Science University, USA
M. Colombetti	Politecnico di Milano, Italy
M. Dastani	Utrecht University, Netherlands
F. Dignum	Utrecht University, Netherlands
R. van Eijk	Utrecht University, Netherlands, (2006)
A. El Fallah-Seghrouchni	University of Paris 6, France
R. Flores	Christopher Newport University, USA, (2006)
F. Guerin	University of Aberdeen, UK
M.-P. Huget	University of Savoie, France, (2006)
M. d’Inverno	Westminster University, UK
A. Jones	King’s College, London, UK
F. Lin	Athabasca University, Canada, (2006)
N. Maudet	University of Paris 9, France
P. McBurney	University of Liverpool, UK
S. Parsons	City University of New York, USA
P. Pasquier	University of Melbourne, Australia, (2006)
S. Paurolally	University of Liverpool, UK
J. Pitt	Imperial College, UK
N. Roos	Maastricht University, Netherlands
D. Traum	University of California Los Angeles, USA
G. Weiss	Technical University Munich, Germany
M. Wooldridge	University of Liverpool, UK, (2006)
P. Yolum	Bogazici University, Turkey

External Reviewers

D. Grossi	Universiteit Utrecht, Netherlands, (2005)
M. Nickles	Technical University of Munich, Germany, (2005)
P. Pasquier	Laval University, Canada, (2005)
M. Rovatsos	University of Edinburgh, UK, (2005)
J. Saunier	University of Paris 9, France, (2006)

Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 4333: U. Reimer, D. Karagiannis (Eds.), *Practical Aspects of Knowledge Management*. XII, 338 pages. 2006.
- Vol. 4304: A. Sattar, B.-H. Kang (Eds.), *AI 2006: Advances in Artificial Intelligence*. XXVII, 1303 pages. 2006.
- Vol. 4293: A. Gelbukh, C.A. Reyes-Garcia (Eds.), *MICA 2006: Advances in Artificial Intelligence*. XXVIII, 1232 pages. 2006.
- Vol. 4289: M. Ackermann, B. Berendt, M. Grobelnik, A. Hotho, D. Mladenić, G. Semeraro, M. Spiliopoulou, G. Stumme, V. Svatek, M. van Someren (Eds.), *Semantics, Web and Mining*. X, 197 pages. 2006.
- Vol. 4285: Y. Matsumoto, R. Sproat, K.-F. Wong, M. Zhang (Eds.), *Computer Processing of Oriental Languages*. XVII, 544 pages. 2006.
- Vol. 4274: Q. Huo, B. Ma, E.-S. Chng, H. Li (Eds.), *Chinese Spoken Language Processing*. XXIV, 805 pages. 2006.
- Vol. 4265: L. Todorovski, N. Lavrač, K.P. Jantke (Eds.), *Discovery Science*. XIV, 384 pages. 2006.
- Vol. 4264: J.L. Balcázar, P.M. Long, F. Stephan (Eds.), *Algorithmic Learning Theory*. XIII, 393 pages. 2006.
- Vol. 4259: S. Greco, Y. Hata, S. Hirano, M. Inuiguchi, S. Miyamoto, H.S. Nguyen, R. Słowiński (Eds.), *Rough Sets and Current Trends in Computing*. XXII, 951 pages. 2006.
- Vol. 4253: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part III*. XXXII, 1301 pages. 2006.
- Vol. 4252: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part II*. XXXIII, 1335 pages. 2006.
- Vol. 4251: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part I*. LXVI, 1297 pages. 2006.
- Vol. 4248: S. Staab, V. Svátek (Eds.), *Managing Knowledge in a World of Networks*. XIV, 400 pages. 2006.
- Vol. 4246: M. Hermann, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning*. XIII, 588 pages. 2006.
- Vol. 4223: L. Wang, L. Jiao, G. Shi, X. Li, J. Liu (Eds.), *Fuzzy Systems and Knowledge Discovery*. XXVIII, 1335 pages. 2006.
- Vol. 4213: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Eds.), *Knowledge Discovery in Databases: PKDD 2006*. XXII, 660 pages. 2006.
- Vol. 4212: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Eds.), *Machine Learning: ECML 2006*. XXIII, 851 pages. 2006.
- Vol. 4211: P. Vogt, Y. Sugita, E. Tuci, C.L. Nehaniv (Eds.), *Symbol Grounding and Beyond*. VIII, 237 pages. 2006.
- Vol. 4203: F. Esposito, Z.W. Raś, D. Malerba, G. Semeraro (Eds.), *Foundations of Intelligent Systems*. XVIII, 767 pages. 2006.
- Vol. 4201: Y. Sakakibara, S. Kobayashi, K. Sato, T. Nishino, E. Tomita (Eds.), *Grammatical Inference: Algorithms and Applications*. XII, 359 pages. 2006.
- Vol. 4200: I.F.C. Smith (Ed.), *Intelligent Computing in Engineering and Architecture*. XIII, 692 pages. 2006.
- Vol. 4198: O. Nasraoui, O. Zaïane, M. Spiliopoulou, B. Mobasher, B. Masand, P.S. Yu (Eds.), *Advances in Web Mining and Web Usage Analysis*. IX, 177 pages. 2006.
- Vol. 4196: K. Fischer, I.J. Timm, E. André, N. Zhong (Eds.), *Multiagent System Technologies*. X, 185 pages. 2006.
- Vol. 4188: P. Sojka, I. Kopeček, K. Pala (Eds.), *Text, Speech and Dialogue*. XV, 721 pages. 2006.
- Vol. 4183: J. Euzenat, J. Domingue (Eds.), *Artificial Intelligence: Methodology, Systems, and Applications*. XIII, 291 pages. 2006.
- Vol. 4180: M. Kohlhase, OMDoc – An Open Markup Format for Mathematical Documents [version 1.2]. XIX, 428 pages. 2006.
- Vol. 4177: R. Marín, E. Onaindía, A. Bugarín, J. Santos (Eds.), *Current Topics in Artificial Intelligence*. XV, 482 pages. 2006.
- Vol. 4160: M. Fisher, W. van der Hoek, B. Konev, A. Lisitsa (Eds.), *Logics in Artificial Intelligence*. XII, 516 pages. 2006.
- Vol. 4155: O. Stock, M. Schaerf (Eds.), *Reasoning, Action and Interaction in AI Theories and Systems*. XVIII, 343 pages. 2006.
- Vol. 4149: M. Klusch, M. Rovatsos, T.R. Payne (Eds.), *Cooperative Information Agents X*. XII, 477 pages. 2006.
- Vol. 4140: J.S. Sichman, H. Coelho, S.O. Rezende (Eds.), *Advances in Artificial Intelligence - IBERAMIA-SBIA 2006*. XXIII, 635 pages. 2006.
- Vol. 4139: T. Salakoski, F. Ginter, S. Pyysalo, T. Pahikkala (Eds.), *Advances in Natural Language Processing*. XVI, 771 pages. 2006.
- Vol. 4133: J. Gratch, M. Young, R. Allett, D. Ballin, P. Olivier (Eds.), *Intelligent Virtual Agents*. XIV, 472 pages. 2006.
- Vol. 4130: U. Furbach, N. Shankar (Eds.), *Automated Reasoning*. XV, 680 pages. 2006.

- Vol. 4120: J. Calmet, T. Ida, D. Wang (Eds.), Artificial Intelligence and Symbolic Computation. XIII, 269 pages. 2006.
- Vol. 4114: D.-S. Huang, K. Li, G.W. Irwin (Eds.), Computational Intelligence, Part II. XXVII, 1337 pages. 2006.
- Vol. 4108: J.M. Borwein, W.M. Farmer (Eds.), Mathematical Knowledge Management. VIII, 295 pages. 2006.
- Vol. 4106: T.R. Roth-Berghofer, M.H. Göker, H.A. Güvenir (Eds.), Advances in Case-Based Reasoning. XIV, 566 pages. 2006.
- Vol. 4099: Q. Yang, G. Webb (Eds.), PRICAI 2006: Trends in Artificial Intelligence. XXVIII, 1263 pages. 2006.
- Vol. 4095: S. Nolfi, G. Baldassarre, R. Calabretta, J.C.T. Hallam, D. Marocco, J.-A. Meyer, O. Miglino, D. Parisi (Eds.), From Animals to Animats 9. XV, 869 pages. 2006.
- Vol. 4093: X. Li, O.R. Zaïane, Z. Li (Eds.), Advanced Data Mining and Applications. XXI, 1110 pages. 2006.
- Vol. 4092: J. Lang, F. Lin, J. Wang (Eds.), Knowledge Science, Engineering and Management. XV, 664 pages. 2006.
- Vol. 4088: Z.-Z. Shi, R. Sadananda (Eds.), Agent Computing and Multi-Agent Systems. XVII, 827 pages. 2006.
- Vol. 4087: F. Schwenker, S. Marinai (Eds.), Artificial Neural Networks in Pattern Recognition. IX, 299 pages. 2006.
- Vol. 4068: H. Schärfe, P. Hitzler, P. Øhrstrøm (Eds.), Conceptual Structures: Inspiration and Application. XI, 455 pages. 2006.
- Vol. 4065: P. Perner (Ed.), Advances in Data Mining. XI, 592 pages. 2006.
- Vol. 4062: G. Wang, J.F. Peters, A. Skowron, Y. Yao (Eds.), Rough Sets and Knowledge Technology. XX, 810 pages. 2006.
- Vol. 4049: S. Parsons, N. Maudet, P. Moraitis, I. Rahwan (Eds.), Argumentation in Multi-Agent Systems. XIV, 313 pages. 2006.
- Vol. 4048: L. Goble, J.-J.C. Meyer (Eds.), Deontic Logic and Artificial Normative Systems. X, 273 pages. 2006.
- Vol. 4045: D. Barker-Plummer, R. Cox, N. Swoboda (Eds.), Diagrammatic Representation and Inference. XII, 301 pages. 2006.
- Vol. 4031: M. Ali, R. Dapoigny (Eds.), Advances in Applied Artificial Intelligence. XXIII, 1353 pages. 2006.
- Vol. 4029: L. Rutkowski, R. Tadeusiewicz, L.A. Zadeh, J.M. Zurada (Eds.), Artificial Intelligence and Soft Computing – ICAISC 2006. XXI, 1235 pages. 2006.
- Vol. 4027: H.L. Larsen, G. Pasi, D. Ortiz-Arroyo, T. Andreassen, H. Christiansen (Eds.), Flexible Query Answering Systems. XVIII, 714 pages. 2006.
- Vol. 4021: E. André, L. Dybkjær, W. Minker, H. Neumann, M. Weber (Eds.), Perception and Interactive Technologies. XI, 217 pages. 2006.
- Vol. 4020: A. Bredenfeld, A. Jacoff, I. Noda, Y. Takahashi (Eds.), RoboCup 2005: Robot Soccer World Cup IX. XVII, 727 pages. 2006.
- Vol. 4013: L. Lamontagne, M. Marchand (Eds.), Advances in Artificial Intelligence. XIII, 564 pages. 2006.
- Vol. 4012: T. Washio, A. Sakurai, K. Nakajima, H. Takeda, S. Tojo, M. Yokoo (Eds.), New Frontiers in Artificial Intelligence. XIII, 484 pages. 2006.
- Vol. 4008: J.C. Augusto, C.D. Nugent (Eds.), Designing Smart Homes. XI, 183 pages. 2006.
- Vol. 4005: G. Lugosi, H.U. Simon (Eds.), Learning Theory. XI, 656 pages. 2006.
- Vol. 4002: A. Yli-Jyrä, L. Karttunen, J. Karhumäki (Eds.), Finite-State Methods and Natural Language Processing. XIV, 312 pages. 2006.
- Vol. 3978: B. Hnich, M. Carlsson, F. Fages, F. Rossi (Eds.), Recent Advances in Constraints. VIII, 179 pages. 2006.
- Vol. 3963: O. Dikenelli, M.-P. Gleizes, A. Ricci (Eds.), Engineering Societies in the Agents World VI. XII, 303 pages. 2006.
- Vol. 3960: R. Vieira, P. Quaresma, M.d.G.V. Nunes, N.J. Mamede, C. Oliveira, M.C. Dias (Eds.), Computational Processing of the Portuguese Language. XII, 274 pages. 2006.
- Vol. 3955: G. Antoniou, G. Potamias, C. Spyropoulos, D. Plexousakis (Eds.), Advances in Artificial Intelligence. XVII, 611 pages. 2006.
- Vol. 3949: F.A. Savacı (Ed.), Artificial Intelligence and Neural Networks. IX, 227 pages. 2006.
- Vol. 3946: T.R. Roth-Berghofer, S. Schulz, D.B. Leake (Eds.), Modeling and Retrieval of Context. XI, 149 pages. 2006.
- Vol. 3944: J. Quiñero-Candela, I. Dagan, B. Magnini, F. d'Alché-Buc (Eds.), Machine Learning Challenges. XIII, 462 pages. 2006.
- Vol. 3937: H. La Poutré, N.M. Sadeh, S. Janson (Eds.), Agent-Mediated Electronic Commerce. X, 227 pages. 2006.
- Vol. 3932: B. Mobasher, O. Nasraoui, B. Liu, B. Masand (Eds.), Advances in Web Mining and Web Usage Analysis. X, 189 pages. 2006.
- Vol. 3930: D.S. Yeung, Z.-Q. Liu, X.-Z. Wang, H. Yan (Eds.), Advances in Machine Learning and Cybernetics. XXI, 1110 pages. 2006.
- Vol. 3918: W.-K. Ng, M. Kitsuregawa, J. Li, K. Chang (Eds.), Advances in Knowledge Discovery and Data Mining. XXIV, 879 pages. 2006.
- Vol. 3913: O. Boissier, J. Padget, V. Dignum, G. Lindemann, E. Matson, S. Ossowski, J.S. Sichman, J. Vázquez-Salceda (Eds.), Coordination, Organizations, Institutions, and Norms in Multi-Agent Systems. XII, 259 pages. 2006.
- Vol. 3910: S.A. Brueckner, G.D.M. Serugendo, D. Hales, F. Zambonelli (Eds.), Engineering Self-Organising Systems. XII, 245 pages. 2006.
- Vol. 3904: M. Baldoni, U. Endriss, A. Omicini, P. Torroni (Eds.), Declarative Agent Languages and Technologies III. XII, 245 pages. 2006.
- Vol. 3900: F. Toni, P. Torroni (Eds.), Computational Logic in Multi-Agent Systems. XVII, 427 pages. 2006.

¥410.00元

Table of Contents

Section I: Semantics of Agent Communication

An Operational Model for the FIPA-ACL Semantics	1
<i>Vincent Louis, Thierry Martinez</i>	
Temporal Logics for Representing Agent Communication Protocols	15
<i>Ulle Endriss</i>	
ACL Semantics Between Social Commitments and Mental Attitudes	30
<i>Guido Boella, Rossana Damiano, Joris Hulstijn, Leendert van der Torre</i>	
On the Semantics of Conditional Commitment	45
<i>Shakil M. Khan, Yves Lespérance</i>	

Section II: Commitments in Agent Communication

A Commitment-Based Communicative Act Library	61
<i>Mario Verdicchio, Marco Colombetti</i>	
Integrating Social Commitment-Based Communication in Cognitive Agent Modeling	76
<i>Philippe Pasquier, Brahim Chaib-draa</i>	
Flexible Conversations Using Social Commitments and a Performatives Hierarchy	93
<i>Rob Kremer, Roberto A. Flores</i>	
Using Social Commitments to Control the Agents' Freedom of Speech . . .	109
<i>Guillaume Muller, Laurent Vercouter</i>	
Practical Issues in Detecting Broken Social Commitments	124
<i>Jason Heard, Rob Kremer</i>	
Introducing Preferences into Commitment Protocols	136
<i>Ashok U. Mallya, Munindar P. Singh</i>	

Section III: Protocols and Strategies

On the Study of Negotiation Strategies	150
<i>Leila Amgoud, Souhila Kaci</i>	
Strategies for Ontology Negotiation: Finding the Right Level of Generality	164
<i>Jurriaan van Diggelen, Edwin D. de Jong, Marco A. Wiering</i>	
Combining Normal Communication with Ontology Alignment	181
<i>Jurriaan van Diggelen, Robbert Jan Beun, Frank Dignum, Rogier M. van Eijk, John-Jules Meyer</i>	
Towards Design Tools for Protocol Development	196
<i>Pinar Yolum</i>	
Adaptiveness in Agent Communication: Application and Adaptation of Conversation Patterns	211
<i>Felix Fischer, Michael Rovatsos, Gerhard Weiss</i>	
Can I Please Drop It? Dialogues About Belief Contraction	227
<i>Henk-Jan Lebbink, Cilia Witteman, John-Jules Meyer</i>	
Commitment-Based Policies in Persuasion Dialogues with Defeasible Beliefs	243
<i>Ioan Alfred Letia, Raluca Vartic</i>	

Section IV: Reliability and Overhearing

Reliable Group Communication and Institutional Action in a Multi-Agent Trading Scenario	258
<i>Stephen Cranefield</i>	
A Fault Tolerant Agent Communication Language for Supporting Web Agent Interaction	273
<i>Nicola Dragoni, Mauro Gaspari, Davide Guidi</i>	
Experiments in Selective Overhearing of Hierarchical Organizations	289
<i>Gery Gutnik, Gal A. Kaminka</i>	
Author Index	303

An Operational Model for the FIPA-ACL Semantics

Vincent Louis and Thierry Martinez

France Telecom Research & Development
2, avenue Pierre Marzin
22307 Lannion Cedex, France

{vincent.louis, thierry.martinez}@francetelecom.com

Abstract. Despite the effort made to standardize agent communication languages, almost no tool has been developed to implement agents' conformance to their semantics. In this paper, we review the formal principles supporting the FIPA-ACL semantics and propose an operational model facilitating their implementation. This model has been implemented upon the JADE platform, resulting in more flexible agents, avoiding intensive use of rigid protocols.

1 Introduction

Many research and industrial actors in the field of multi-agent systems have identified the need for a shared agent communication language (ACL) long ago. The most enthusiastic ones consider such a language as the counterpart of human natural languages for agents. In particular, ACLs should make it possible to convey meanings instead of "simple" objects with no semantics like in classical object middlewares [1]. At least, ACLs should let heterogeneous agents communicate and interact with each other [2]. This trend resulted in late nineties in mainly two initiatives to come to a standard language: KQML (Knowledge Query and Manipulation Language) from the ARPA knowledge sharing project [3] and FIPA-ACL from the Foundation for Intelligent Physical Agents consortium [4, 5]. Although the usefulness of such languages in building open, heterogeneous and interoperable agent systems is generally acknowledged, ACLs have also often been criticized because of their formal semantics, which make them generally difficult to implement and hence seldom implemented [6, 7].

On the one hand, a commonly pointed out difficulty in using these ACLs is that they assume a mental state-based model for agents. While such mentalistic approaches are well suited to specify the meaning of communicative acts from the subjective viewpoint of the participants, they cannot ensure any global objective property of the system. Along this line, a recent stream in the multi-agent community has proposed to use so-called social approaches (as opposed to mentalistic or individual ones) [2]. These approaches model interactions through public structures, generally based on the notion of (social) commitment. They lead to recent concrete alternatives for defining ACLs [8, 9].

On the other hand, because the formal semantics of ACLs rely on complex logical theories of agency (mixing reasoning on several concepts such as beliefs, intentions and actions), they require a high level of understanding to design conforming agents. This matter of fact is not only pointed out at a theoretical level but also at a concrete

development level [10]. Consequently, even if some recent work attempts to simplify this kind of models [11], designers often prefer protocol-based approaches (such as [12]). In these cases, they do not benefit from the original semantics of ACLs, which only account for an intuitive meaning of the communicative acts. As a resulting drawback, the use of rigid interaction protocols often results in decreasing the flexibility and therefore the autonomy of agents.

A third possible working direction consists in making available proper tools in order to encompass these difficulties. To our knowledge, little work has investigated this area. For example, in the FIPA community, most of the existing platforms that claim to be FIPA-compliant (among the most famous ones, JADE [13], FIPA-OS [14], Zeus [15]) implement the middleware-related specifications but provide no concrete support regarding the ACL semantics-related specifications. In order to promote such tools, this paper proposes an operational model for implementing the theory of agency that underlies the FIPA-ACL semantics. This model, which is obviously missing today, aims at both helping developers to soundly conform to the FIPA-ACL semantics and leading to the development of proper tools to be integrated into FIPA-compliant platforms. Actually, this model provides a design framework that ensures consistency with the theory principles. Interestingly, it is flexible enough to customize agents' behaviors, while the built agents also benefit from generic capabilities for interpreting and generating communicative acts. This paper only focuses on the interpretation part.

The next section reviews some formal principles of the theory of agency underlying the FIPA-ACL semantics, which are relevant to the interpretation of communicative acts. Section 3 describes the main concepts and mechanisms of the model from which these principles can be implemented. Section 4 illustrates the resulting model with a simple example. Finally, section 5 concludes and discusses some perspectives.

2 Reviewing FIPA-ACL Semantics

The FIPA Agent Communication Language is defined through a set of communicative acts [5]. Their precise meaning results from their interpretation as particular actions within a more general theory of agency, namely the theory of rational interaction proposed by Sadek [16]. Thus, the semantics of FIPA communicative acts is formally defined by the generic principles of Sadek's theory that apply to actions. Although the FIPA specifications list most of these principles [5, Informative Annex A], some significant ones are unfortunately missing. This section reviews the essential principles formalizing the interpretation of FIPA-ACL communicative acts and identifies a general template in the perspective of implementing them.

All formal properties described in this paper are written in FIPA-SL [17], which is the modal logic language that sustains the theory of agency defining the FIPA-ACL semantics.

2.1 Mental Attitudes and Actions

First of all, the underlying theory formally specifies agents' behaviors through mental state notions describing internal agents' features that must be interpreted subjectively,

i.e. from their point of view. Mental states are classically described using beliefs and intentions (according to the widely acknowledged Belief-Desire-Intention paradigm [18]). Beliefs are formalized by two logical modal operators: $(B \underline{i} \underline{p})$ ¹ expresses that agent \underline{i} believes that \underline{p} holds and $(U \underline{i} \underline{p})$ expresses that agent \underline{i} is uncertain about \underline{p} , that is, s/he does not believe that \underline{p} holds but s/he believes that \underline{p} is more probably true than $(\text{not } \underline{p})$. The B operator satisfies a KD45 model and is semantically defined by a Kripke possible world structure, whereas the semantics of the U operator is defined in probabilistic terms upon the accessible possible worlds supporting the B operator.² Intentions are mainly formalized by one logical operator, $(I \underline{i} \underline{p})$, which expresses that agent \underline{i} intends that \underline{p} holds. Sadek's theory actually provides several degrees of intention (similar, to some extent, to Cohen and Levesque's approach [20]), including choice, achievement goal (agent \underline{i} does not believe that \underline{p} holds), persistent goal (agent \underline{i} will drop her/his goal \underline{p} until it is satisfied or s/he comes to believe it is unachievable) and intention itself (agent \underline{i} commits to perform, individually or collectively with other cooperative agents, any action s/he believes that can reach the goal \underline{p}), each one being defined upon a more primitive choice concept [21]. Within the scope of the FIPA specifications, the intention operator is considered to be primitive.

A property of agents' mental attitudes, which is worth mentioning, is that they must be consistent with their beliefs: agents always believe the mental attitudes they actually have. Formally, the following property is valid within the theory, for both primitive (expressed with the previous operators of belief, uncertainty and intention) and composite mental attitudes (expressed by combining these operators with logical connectors):

$$(\text{equiv } (B \underline{i} \underline{\text{PHI}(\underline{i})}) \underline{\text{PHI}(\underline{i})}) \quad (1)$$

where \underline{i} denotes an agent and $\underline{\text{PHI}(\underline{i})}$ a mental attitude of \underline{i}

Consequently, agents' internal states (including all their mental attitudes) can be exclusively described by their beliefs. An important corollary is that agents cannot be uncertain (with the meaning of the U operator) of any of their mental attitudes (since they fully believe it). The following property is valid within the theory:

$$(\text{not } (U \underline{i} \underline{\text{PHI}(\underline{i})})) \quad (2)$$

where \underline{i} denotes an agent and $\underline{\text{PHI}(\underline{i})}$ a mental attitude of \underline{i}

In order to describe temporal facts, the theory supports two other modalities: $(\text{done } \underline{a} \underline{p})$ expresses that action \underline{a} has just occurred and that \underline{p} held just before its occurrence (past-oriented), and $(\text{feasible } \underline{a} \underline{p})$ expresses that action \underline{a} may possibly occur and that \underline{p} will hold just after its occurrence, if it actually occurs (future-oriented). Both these operators are possible normal modal operators that satisfy a K model and are semantically defined by a Kripke possible world structure. Their accessibility relations classically define a branching future (several different actions

¹ In this paper, the underlined terms in logical formulas denote schematic variables. Here, \underline{i} and \underline{p} may be respectively replaced with references to agents and formulas.

² Sadek originally proposed a set of logical properties satisfied by the U operator. More recent work investigates an axiomatic system for this operator, but without proving its completeness with respect to the semantic model [19].

may occur in a given possible world) and a linear past (exactly one action has just occurred in a given possible world).

The resulting framework is a homogeneous multimodal logic powerful enough to account for very subtle nuances. For example, $(\text{exists } ?X \ (B \ i \ (\text{feasible } ?X \ p)))$ expresses that agent i knows an action $?X$ (which is not explicit) that may bring about p (e.g. “Mary knows a recipe to cook a cake”), whereas $(B \ i \ (\text{exists } ?X \ (\text{feasible } ?X \ p)))$ expresses that agent i believes that there are some ways of reaching p , without necessarily knowing how to reach it (e.g. “Mary knows cakes can be cooked”).

2.2 Formally Interpreting Incoming Communicative Acts

FIPA-ACL defines four primitive (Inform, Confirm, Disconfirm and Request) and eighteen composite communicative acts [5]. Each of them (primitive or composite) is defined by two semantic features, namely its feasibility precondition and its rational effect. We now review the formal principles embedded in the theory of agency that specify how agents should interpret received communicative acts based on their semantic features.

Feasibility Precondition

The feasibility precondition of a communicative act states the condition that must necessarily hold for this act to be sent. This classical notion of action precondition is formalized by the following axiom within the theory³. It means that any agent observing a communicative act performance (left part of the implication) necessarily believes that its feasibility precondition held just before its performance (right part, in the scope of the done operator):

$$(B \ i \ (\text{implies} \ (\text{done } \underline{a} \ \text{true}) \ (\text{done } \underline{a} \ \underline{FP}(\underline{a})))) \quad (3)$$

where \underline{a} , $\underline{FP}(\underline{a})$ and \underline{i} respectively denote a communicative act, its feasibility precondition and an agent

This principle is particularly useful to check the consistency of incoming communicative acts. For example, agents should reject received Inform acts about one of their own mental attitudes (e.g. when they are told “you intend to jump out the window”) because applying property (2) to the corresponding propositional content makes the feasibility precondition inconsistent. Actually, the informative annex of [5] mentions no property that formally deals with feasibility precondition interpretation, so that a specification of inconsistent communicative acts is clearly missing.

Rational Effect

The rational effect of a communicative act states what the result expected by agents performing this act is. It underlies a unique classical actual postcondition of the communicative act, namely its “intentional effect”, which is formalized by the following axiom within the theory. It means that agents observing an act performance (left part of the implication) believe that the sender intends each receiver believes the sender intends the rational effect of this act (right part):

³ This axiom, which was actually proposed by Louis [19], generalizes Sadek’s original formalization: $(B \ i \ (\text{implies} \ (\text{feasible } \underline{a} \ \text{true}) \ \underline{FP}(\underline{a}))))$.