

08361 Abstracts Collection
Programming Multi-Agent Systems
— **Dagstuhl Seminar** —

Rafael Bordini¹, Mehdi Dastani², Jürgen Dix³ and Amal El
Fallah-Seghrouchni⁴

¹ University of Durham, United Kingdom

`r.bordini@durham.ac.uk`

² Utrecht University, Netherlands

`mehdi@cs.uu.nl`

³ TU Clausthal, Germany

`dix@tu-clausthal.de`

⁴ Université Paris VI, France

`Amal.Elfallah@lip6.fr`

Abstract. From 31th August to 5th September, the Dagstuhl Seminar 08361 “Programming Multi-Agent Systems” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Agency, agent-oriented programming, programming multi-agent systems

08361 Executive Summary – Programming Multi-Agent Systems

The term agent-oriented programming language was coined by Yoav Shoham (Stanford University) almost 20 years ago. He presented both a new concept, and a prototype programming language (AGENT0), which embodied some important ideas and led to a new programming paradigm. Since then, these ideas have been improved, refined, and extended by many researchers and led to a whole new subfield of computer science. The importance of MAS is also reflected in the number of conferences, workshops and journals that have evolved over the last ten years. We are convinced that the next step in furthering the achievement of the multiagent systems project is irrevocably associated with the development of programming languages and tools that can effectively support

2 Rafael Bordini, Mehdi Dastani, Jürgen Dix and Amal El Fallah-Seghrouchni

multi-agent programming, including the implementation of key notions in multi-agent systems in a unified framework. We hope this seminar will pave the way to many useful contributions in that direction.

We aim to bring together a combination of academics and industrialists with diverse expertise. Because of the nature of this research area, multi-agent system programming has attracted researchers who come from various backgrounds, all of which have fundamental contributions to make to this huge endeavour.

Keywords: Multi-Agent Programming, Multi-Agent Systems

Joint work of: Bordini, Rafael; Dastani, mehdi; Dix, Jürgen; El Fallah-Seghrouchni, Amal

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2008/1646>

A Rewrite Based Framework for Designing and Verifying BDI Agents

Lacramioara Astefanoaei (CWI - Amsterdam, Netherlands)

Rewriting logic has the benefit that there is no gap between semantics and implementation; it allows rapid prototyping, and it offers a suite of generic tools for verification.

In this talk, the focus is on rewriting logic as a framework for incrementally designing modular agent languages and for formally analysing their definitions. Some experiments with model-checking BDI-like agents in Maude are also presented.

Keywords: BDI, Non-determinism, Refinement, Rewriting, Simulation

Model-Checking Agent Refinement

Lacramioara Astefanoaei (CWI - Amsterdam, Netherlands)

We present a proof-technique for reducing the nondeterminism of abstract agent specifications in a BDI framework by means of refinement. We implement the operational semantics of agent specifications in rewrite systems such that we can automatically check if refinement between (fair) executions of agents holds.

Keywords: BDI, Non-determinism, Refinement, Rewriting, Simulation

Joint work of: Astefanoaei, Lacramioara; de Boer, Frank S.

Full Paper:

<http://homepages.cwi.nl/~astefano/model-check-agent-refinement.pdf>

See also: F. Arbab. Reo: a channel-based coordination model for component composition. *Mathematical Structures in Computer Science*, 14(3):329–366, 2004.

R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. Verifying multi-agent programs by model checking. *Autonomous Agents and Multi-Agent Systems*, 12(2):239–256, 2006.

T. Bosse, C. M. Jonker, L. van der Meij, A. Sharpanskykh, and J. Treur. Specification and verification of dynamics in cognitive agent models. In *IAT*, pages 247–254, 2006.

M. Bratman. *Intentions, Plans, and Practical Reason*. Harvard University Press, 1987.

K. M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison Wesley Publishing Company, Inc., Reading, Massachusetts, 1988.

M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and J. F. Quesada. *Maude: Specification and programming in rewriting logic*. *Theoretical Computer Science*, 2001.

M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. L. Talcott. *All About Maude - A High-Performance Logical Framework*, volume 4350 of *Lecture Notes in Computer Science*. Springer, 2007.

F. S. de Boer, K. V. Hindriks, W. van der Hoek, and J.-J. C. Meyer. A verification framework for agent programming with declarative goals. *J. Applied Logic*, 5(2):277–302, 2007.

J. Edmund M. Clarke, O. Grumberg, and D. A. Peled. *Model checking*. MIT Press, Cambridge, MA, USA, 1999.

S. Eker, J. Meseguer, and A. Sridharanarayanan. The Maude LTL model checker and its implementation. In *Model Checking Software: Proc. 10 th Intl. SPIN Workshop*, volume 2648 of *LNCS*, pages 230–234. Springer, 2003.

K. V. Hindriks, F. S. de Boer, W. van der Hoek, and J.-J. C. Meyer. Agent programming in 3apl. *Autonomous Agents and Multi-Agent Systems*, 2(4):357–401, 1999.

Z. Manna and A. Pnueli. *The temporal logic of reactive and concurrent systems*. Springer-Verlag New York, Inc., New York, NY, USA, 1992.

J. Meseguer and G. Rosu. *Rewriting logic semantics: From language specifications to formal analysis tools*.

F. Raimondi and A. Lomuscio. Automatic verification of multi-agent systems by model checking via ordered binary decision diagrams. *J. Applied Logic*, 5(2):235–251, 2007.

S. Safra. *Complexity of automata on infinite objects*. PhD thesis, Rehovot, Israel, 1989.

M. B. van Riemsdijk, F. S. de Boer, M. Dastani, and J.-J. C. Meyer. Prototyping 3apl in the maude term rewriting language. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 1279–1281, New York, NY, USA, 2006. ACM Press.

A. Verdejo and N. Martí-Oliet. Two case studies of semantics execution in maude: Ccs and lotos. *Form. Methods Syst. Des.*, 27(1/2):113–172, 2005.

Choice, Interoperability, and Conformance in Service Choreographies

Matteo Baldoni (University of Torino, Italy)

The accomplishment of a complex task often requires interaction among a set of interoperable services. At every point of such an interaction, one of the services has the (implicit) right of deciding which action to take, while its partner must have the ability to tackle every such decision.

This paper proposes notions of interoperability and conformance that are based on an explicit representation of which partner chooses at each step and which partner respects the choice. Using these notions, this paper shows how to approach (the previously unaddressed) multiparty interactions. It also introduces a set of edit operations that enable the construction of conformant variants of a service implementation as well as of a role specification in a service choreography. These operations enable the construction of compatible service upgrades or patches that make a service implementation interoperable with an interlocutor with which it was not originally compatible.

Keywords: Protocol verification, conformance, communication

Joint work of: Baldoni, Matteo; Baroglio, Cristina; Chopra, Amit K.; Desai, Nirmitt; Patti, Viviana; Singh, Munindar P.

Choreography-driven matchmaking and selection of web services

Cristina Baroglio (University of Torino, Italy)

The greater and greater quantity of services that are available over the web causes a growing attention to techniques that facilitate their reuse. A web service specification can be quite complex, including various operations and message exchange patterns. In our recent work, we have focussed on the problem of selecting a web service, which can play a role in a given choreography, preserving at the same time a condition of interest (the goal that caused the search of the service). In particular, after proving that current semantic matchmaking techniques do not guarantee goal preservation, we are completing an approach for overcoming these limits. This approach exploits the choreography definition and is based on an action-based representation of service operations: each operation is described in terms of its preconditions and effects, without taking into account the ontology layer, which is not functional to the aims of the work.

Keywords: Matchmaking, reasoning about conversation, verification

Joint work of: Baldoni, Matteo; Baroglio, Cristina; Martelli, Alberto; Patti, Viviana; Schifanella, Claudio

Agent Contest Competition: 4th Edition

Tristan M. Behrens (TU Clausthal, Germany)

The aim of the contest is to stimulate research in the area of multi-agent systems by identifying key problems and collecting suitable benchmarks that can serve as milestones for evaluating new tools, models, and techniques to develop multi-agent systems. The first two editions of this contest were organised in association with CLIMA conference series and the third edition was organised in association with ProMAS'07. Based on the experiences from the previous three editions, the contest scenario has been changed to test the participating multi-agent systems on their abilities to coordinate and cooperate. We wanted to emphasise team work and team strategy issues in a dynamic environment where teams compete for the same resources. Seven groups from Iran, Ireland, England, France, Germany, Poland, and Turkey did participate in this years contest.

Keywords: Multi-agent systems programming, benchmarks, comparison

Joint work of: Behrens, T. M. ; Dastani, Mehdi; Dix, Jürgen; Novak, P.

ORA4MAS: an organisational environment for MAS

Olivier Boissier (Ecole des Mines - St. Etienne, France)

The social and organisational aspects of agency have become nowadays a major focus of interest in the MAS community, and a good amount of theoretical work is available, in terms of formal models and theories. However, the conception and engineering of proper organisational infrastructures embodying such models and theories is still an open issue, in particular when open MAS are considered. In this talk we discuss a model for an organisational infrastructure called ORA4MAS that aims at addressing these issues. By being based on the A&A (Agents and Artifacts) meta-model, the key and novel aspect introduced with ORA4MAS is that organisations and the organisation infrastructure itself are conceived in terms of agents and artifacts, as first-class abstractions giving body to the MAS from design to runtime.

Keywords: MAS organisation, artifacts, MOISE

Joint work of: Boissier, Olivier; Kitio, Rosine; Hubner, Jomi F; Ricci, Alessandro

Instrumenting Multi-Agent Organisation with Organisational Artifacts

Olivier Boissier (Ecole des Mines - St. Etienne, France)

The social and organisational aspects of agency have become nowadays a major focus of interest in the MAS community, and a good amount of theoretical work is available, in terms of formal models and theories.

However, the conception and engineering of proper organisational infrastructures embodying such models and theories is still an open issue, in particular when open MAS are considered. In this talk we discuss a model for an organisational infrastructure called ORA4MAS that aims at addressing these issues. By being based on the A&A (Agents and Artifacts) meta-model, the key and novel aspect introduced with ORA4MAS is that organisations and the organisation infrastructure itself are conceived in terms of agents and artifacts, as first-class abstractions giving body to the MAS from design to runtime.

Keywords: Organisation, Artifact, Organisation Oriented Programming

Joint work of: Boissier, Olivier; Hübner, Jomi; Kitio, Rosine; Ricci, Alessandro

Rule-based agent architectures: Jadex V2

Lars Braubach (Universität Hamburg, Germany)

In a recent comparison of several agent platforms it has been found that one outstanding feature of commercial agent platforms is their excellent operating ability. The overall operating ability has been assessed by evaluating the sub-criteria performance, stability, robustness and scalability of the different representatives. In order to improve the operating ability of the Jadex agent platform, a new BDI agent interpreter architecture called Jadex V2 has been proposed. This architecture is completely based on proven rule engine technology. The forward chaining rule engine is thereby based on a Rete pattern matcher and an object-attribute-value (OAV) state representation. The BDI functionality has been realized as a separate layer on top of the rule engine layer by introducing a set of BDI reasoning rules. These rules implement the whole BDI reasoning process including the goal deliberation and means-end reasoning phases. Despite this fundamentally changed architecture one important objective consisted in preserving the intuitive Jadex V1 programming model. For this purpose, especially the procedural access to the OAV state elements has been shielded by a flyweight layer, which mediates between the internal (OAV-based) and external (Java-based) access facilities.

Keywords: Agent platforms, BDI agent model

Joint work of: Braubach, Lars; Pokahr, Alexander

Reasoning About Rational Agents and Coalitions Using ATL

Nils Bulling (TU Clausthal, Germany)

Alternating-time temporal logic (ATL) is a temporal logic that can be used to model and to reason about multi-agent systems.

The logic incorporates some basic game theoretic notions and has already been extended by additional concepts by several researchers, including ourselves.

In this talk, I will give an overview of three extensions of ATL which focus on rational behavior in multi-agent systems. The first proposed logic, ATL with plausibility (ATLP), allows to focus on rationally acting agents; the underlying assumption is that agents are often not completely dumb and behave according to some rationality assumption. Then, I will speak about Coalitional ATL, a combination of ATL with an argumentation-based approach to coalition formation.

Finally, our most recent extension, ATL with probabilistic success, is presented. This logic adds a probabilistic dimension to ATL in order to predict agents' behavior and tries to soften the rigorous notion of success that underpins ATL.

Keywords: Modal Logic, Multi-agent systems, Rational Agents, Game theory

AgentMT(TR): a multi-threaded robot control architecture using Teleo-Reactive Plans

Keith L. Clark (Imperial College London, United Kingdom)

In this talk we will argue that a multi-threaded control architecture with a library of partial plans that are a generalization of Nilsson's Teleo-Reactive (TR) rule procedures allows for the smooth integration of three key levels of robot control:

1. Speedy but goal directed response to changing sensor readings
2. Switching between level 1 control procedures when higher level inferred beliefs about the state of the environment change
3. Reacting to events and goals by selecting appropriate level 2 control plans

We start by introducing and demonstrating the use of TR rule procedures with a simulation and visualisation of purely reactive can collecting robot.

We will then show how TR procedures, with rules that test beliefs inferred from sensor reading using concept defining rules, can be used to program a block stacking robot arm that can systematically build a tower of named blocks and 'knows' when the tower is completed. A simulation will be demonstrated.

A key feature of TR procedure control is that the robot can be helped or hindered in its task and the TR procedure will immediately respond by skipping actions, if helped, or by redoing actions if hindered. This operational semantics leads naturally to a multi-threaded implementation.

A multi-tasking robot can respond to events: new goal events or just significant belief updates triggered by sensor readings. It then selects an appropriate plan of action for each event using event/plan selection rules.

We conclude by describing our top level control architecture which borrows from AgentSpeak(L). It has an event handling thread that selects and invokes appropriate plans for events, sometimes suspending currently executing but incompatible plans for lower priority events. Its plans are a generalization of TR procedures.

Normative Multi-Agent Programs and Their Logics

Mehdi Dastani (Utrecht University, Netherlands)

Multi-agent systems are viewed as consisting of individual agents whose behaviors are regulated by an organization artefact. This paper presents a simplified version of a programming language that is designed to implement norm-based artefacts. Such artefacts are specified in terms of norms being enforced by monitoring, regimenting and sanctioning mechanisms. The syntax and operational semantics of the programming language are introduced and discussed. A logic is presented that can be used to specify and verify properties of programs developed in this language.

Keywords: Multi-agent systems, norms, agent-oriented programming

Joint work of: Dastani, Mehdi; Grossi, Davide; Meyer, Jophn-Jules Ch.; Tinemeier, Nick

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1637>

Automated Verification of Multi-Agent Programs

Louise Dennis (University of Liverpool, United Kingdom)

In this paper, we show that the flexible model-checking of multi-agent systems, implemented using agent-oriented programming languages, is viable thus paving the way for the construction of verifiably correct applications of autonomous agents and multi-agent systems.

Model checking experiments were carried out on AJPF (Agent JPF), our extension of Java PathFinder that incorporates the Agent Infrastructure Layer, our unifying framework for agent programming languages. In our approach, properties are specified in a temporal language extended with (shallow) agent-related modalities. The framework then allows the verification of programs written in a variety of agent programming languages, thus removing the need for individual languages to implement their own verification framework.

It even allows the verification of multi-agent systems comprised of agents developed in a variety of different (agent) programming languages.

Keywords: Model checking, BDI

Task Suspension in Agent Systems

Berndt Farwer (University of Durham, United Kingdom)

We discuss the similarity of a recent approach to task suspension in agent programming languages with an earlier approach to formalising preemption using a class of Petri nets, called M-nets.

We argue that the theory of agent programming would benefit from adopting certain features of the Petri-net approach, and thus making further results for Petri nets applicable in the agent domain.

Keywords: Agent programming, Petri nets, suspension, preemption

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1638>

Formal Development of Multi-Agent Systems

Michael Fisher (University of Liverpool, United Kingdom)

In this talk I will describe some of our ongoing work on providing a formal development framework for dynamic multi-agent systems. Although this is based on executable agent specifications, it also explores the development and analysis of organizations.

Keywords: Multi-agent systems; programming; specification; context

Joint work of: Fisher, Michael; Dennis, Louise; Hepple, Anthony

Adding Organizations and Roles as Primitives to the JADE Framework

Roberto Grenna (University of Torino, Italy)

The organization metaphor is often used in the design and implementation of multiagent systems. However, few agent programming languages provide facilities to define them. Several frameworks are proposed to coordinate MAS with organizations, but they are not programmable with general purpose languages. In this paper we extend the JADE framework with primitives to program in Java organizations structured in roles and to enable agents to play roles in organizations. Roles facilitate the coordination of agents inside an organization and give new abilities in the context of organizations, called powers, to the agents which satisfy the requirements necessary to play the roles. As primitives to program organizations and roles we provide classes and protocols which enable an agent to enact a new role in an organization and to interact with the role by invoking the execution of powers, and to receive new goals to be fulfilled. Roles have state and behaviour, thus, they are instances of classes and are strictly connected with the organization offering them. Since roles and organizations can be on a different platform with respect to the role player, the communication with them happens via protocols. Moreover, since, besides using protocols, roles and organizations can have complex behaviours, they are implemented by extending the JADE agent class.

Keywords: Roles, Organizations, JADE

10 Rafael Bordini, Mehdi Dastani, Jürgen Dix and Amal El Fallah-Seghrouchni

Joint work of: Grenna, Roberto; Baldoni, Matteo; Boella, Guido; van der Torre, Leendert; Dorni, Mauro; Mugnaini, Andrea; Genovese, Valerio

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1639>

GOAL Agents Instantiate Intention Logic

Koen V. Hindriks (Delft University of Technologies, Netherlands)

It is commonly believed there is a big gap between agent logics and computational agent frameworks. We show that this gap is not as big as believed by showing that GOAL agents instantiate Intention Logic of Cohen and Levesque. That is, we show that GOAL agent programs can be formally related to Intention Logic. We do so by proving that the GOAL Verification Logic can be embedded into Intention Logic. It follows that (a fragment of) Intention Logic can be used to prove properties of GOAL agents. The work reported is an important step towards the application of standard tools from modal logic for e.g. model checking agent programs. Our results also prove useful for extending the expressiveness of the GOAL agent language. This is illustrated by incorporating temporally extended goals into GOAL agents.

Keywords: Logic, agent program, beliefs, goals

Joint work of: Hindriks, Koen V.; van der Hoek, Wiebe

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1645>

Full Paper:

http://mmi.tudelft.nl/index.php?option=com_pub&task=my_show&Itemid=47

Jason and MOISE+: Organisational Programming in the Agent Contest 2008

Jomi Fred Hübner (Ecole des Mines - St. Etienne, France)

In this talk we report our experience on using organisational programming to develop a team of agents to participate in the Agent Contest 2008 – the cowboys scenario. The main goal is to join BDI like programming in the agent level and the organisational specification and coordination in the system level so that the team may be designed in a high abstraction level (using concepts like groups and roles). Jason was used to code the agents and MOISE+ to specify the organisation. Although these two tools were already integrated, issues like organisational programming style and debugging are still open and the focus of the talk.

Keywords: Jason, MOISE, MAS organisation

Joint work of: Hübner, Jomi F.; Bordini, Rafael H; Picard, Gauthier

Programming Emergent Agents with Multi-Valued Logics

Wojtek Jamroga (TU Clausthal, Germany)

While the ProMAS community has come up with a number of logic-based languages for programming agents, it produced few systems that cope with real (or at least realistic) scenarios.

In consequence, one may suspect that logic is not necessarily well suited for the task. I claim that this is partially due to the fact that logic-based methods provide too coarse characteristics of the domain that is modeled, which makes emergence of behavior very hard to obtain. The reason lies perhaps in the two-valued evaluation of behaviors, offered by classical logic; in particular, gradient search (for a reasonably good behavior) is not possible because the gradient in the space of solutions is either 0 or infinity. For the same reason, emergence is much better obtained by methods that use quantitative modeling of the environment, e.g., soft computing or reinforcement learning.

On the other hand, quantitative methods assume a relatively simple mental structure of agents (e.g., very simple structure of objectives), which can hamper robustness of agents in more complex domains. In contrast, logic-based agents may prove quite robust in the long run... if they can deal with simple tasks (which they can't so far). In this talk, I propose that real-valued logics may be a good answer to the dilemma, by combining logic-like specifications with quantitative modeling and analysis. Finally, I present my recent idea of Markov Temporal Logic (MTL), and unabashedly speculate on how it is going to save the field of ProMAS.

Keywords: Temporal logic, Markov decision process, multi-valued logic

Designing Goal-Oriented Reactive Behaviours

Michael Koester (TU Clausthal, Germany)

In this talk we present some design considerations we faced while implementing Jazzbot, a softbot embodied in a simulated 3D environment, implemented in a rule-based framework of Behavioural State Machines. Finally, we also make an attempt to lift our experiences to a set of informal design guidelines useful for design and implementation of agents with heterogeneous knowledge bases in rule-based agent oriented programming languages.

Keywords: State-of-the-art rule-based agent, jazzbot, simulated 3D environment, behavioural state machines, design guidelines, heterogeneous knowledge bases, answer set programming

Joint work of: Koester, Michael; Novak, Peter

Full Paper:

http://www.aronde.net/uploads/tx_pubdb/GOB_CogRob2008.pdf

Non-Markovian Agent Evolution with EVOLP

João A. Leite (Universidade Nova de Lisboa, Portugal)

Logic Programming Update Languages were proposed as an extension of logic programming, which allow for modelling the dynamics of knowledge bases where both extensional knowledge (facts) as well as intentional knowledge (rules) may change over time due to updates, with important application Multi-Agent Systems (MAS).

Despite their generality, these languages do not provide means to directly access past states of the evolving knowledge. They only allow for so-called Markovian changes i.e. changes determined entirely by the current state. This is a drawback in several situation.

In this paper, after motivating the need for non-Markovian changes, we extend EVOLP – The Logic Programming Update Language at the heart of an existing MAS – with LTL-like temporal operators that allow referring to the history of the evolving agent. We then show that with a suitable introduction of new propositional variables it is possible to embed the extended EVOLP into the original one, thus demonstrating that EVOLP itself can already be used for non-Markovian changes. While showing how to use EVOLP for encoding non-Markovian changes, this embedding sheds light into the relationship between Logic Programming Update Languages and Modal Temporal Logics, of particular importance in MAS.

Keywords: EVOLP, Evolution, Logic Programming

Joint work of: Alferes, José J.; Gabaldon, Alfredo; Leite, João A.

Probabilistic reasoning for agents

John W. Lloyd (Australian National University - Canberra, Australia)

Action selection by agents involves a number of considerations an important one of which is coping with uncertainty. Agents must have some way of modelling uncertainty in their beliefs and some way of reasoning with these uncertain beliefs to select an appropriate action. Looked at from a more general perspective, the problem is a special case of the long-standing one of integrating logic and probability.

In this talk, I will use quantified modal logic to give an elegant solution to this problem. I will show how this logic can capture beliefs that are modal and probabilistic, describe a computation system that an agent can use to reason about such beliefs, and give several examples to illustrate the ideas.

Keywords: Agents, uncertainty, probabilistic beliefs, probabilistic reasoning

Using Goal Decomposition Trees to specify and verify agents and multi-agent systems

Bruno Mermet (University of Le Havre, France)

During this talk, I will introduce the Goal Decomposition Tree (GDT) model that allows to formally specify agents and multiagent systems. The key notion of the model will be presented: our typology of goals, the satisfaction conditions (SC) and the Guaranteed properties in case of failure (GPF), the operators, the proof schemas and the automata composition patterns. I will also present some interesting characteristics of GDTs on a methodological point of view

Keywords: GDT, multiagent systems, formal specification, methodology

Argumentation for Flexible and Adaptive Agent Reasoning

Sanjay Modgil (King's College - London, United Kingdom)

A key challenge for agent architectures and programming paradigms is to account for agent defeasible reasoning and conflict resolution.

A growing body of work is looking to apply argumentation to address this challenge. Much of this work conforms to Dung's seminal abstract argumentation semantics. In this talk I review recent work in which Dung's semantics has been extended to integrate arguments that express meta-level preferences between other arguments. In this way one can account for the fact that preference information required to resolve conflicts is itself defeasible and may be conflicting. I propose the extended argumentation semantics as a semantics for flexible and adaptive agent defeasible and practical reasoning, and show how the extended semantics can characterise, and provide a framework for extending, existing approaches to agent reasoning over beliefs, goals, and actions. For example, agent behavioural types in the *BOID* architecture equate with static preference policies over conflicting goal types. However, agents will need to engage in meta-level reasoning and conflict resolution over these preference policies in order to exhibit behavioural heterogeneity and adapt to different contexts.

Keywords: Argumentation, defeasible reasoning, conflict resolution, preferences

Towards pragmatics of rule-based agent programming language(s)

Peter Novak (TU Clausthal, Germany)

In my research project, I am developing Jazzyk, a special purpose programming language for cognitive agents, with a special stress towards applications in cognitive robotics, such as e.g. small entertainment robots.

Beyond a very generic syntax and semantics, it allows employment of heterogeneous knowledge representation technologies and provides means for a management of interleaving active behaviours of an agent.

In my talk I approach pragmatics of programming Jazzyk, however some of the ideas and results are transferable also to other rule-based languages, such as Jason, GOAL, or 2APL. My main interest is in the process of decomposing a formal specification into a concrete implemented program in a process of a sequel of gradual refinement steps. In order to assist a software designer in this task, I propose a use of design patterns, code templates and common idioms expressed as semantically clearly characterized macros. To characterize the semantics of such code templates, I propose a use of a mixture of Harel's Dynamic Logic together with Temporal Logics, in particular the use of Dynamic Linear Time Temporal Logic.

Keywords: Agent oriented programming languages; temporal logics, DLTTL, goals, dynamic logic, verification

ReSpecT: Teaching Old Dog New Tricks

Andrea Omicini (University of Bologna, Italy)

The agent paradigm smoothly works as a conceptual framework for re-casting old, useful and reliable technologies from many different sources. As a case study, in this talk we discuss how a well-established programming language for the coordination of concurrent process – namely, ReSpecT – can be fully re-interpreted and re-designed by adopting a no-compromise agent-oriented viewpoint. Then, we show how its extension leads to a full-fledged agent-oriented programming language, as well as to a complete and well-rounded language for the coordination of concurrent and distributed processes.

Keywords: MAS coordination, programming languages, ReSpecT

Goal Directed Agents using Artifacts

Michele Piunti (ISTC - CNR - Roma, Italy)

Taking the perspective of Ecosystems and human societies, where living and non-living entities are meant to interact in the same complex environment, we introduce the notion of artifacts in MAS. In particular, we investigate the use of a special kind of artifacts, cognitive artifacts, as computational entities designed to be read and exploited by cognitive agents in order to achieve their goals.

By adopting a functional approach, cognitive artifacts are supposed to play pivotal roles for agents capable to reason about them. On the one side, artifact representational function allows agent to improve epistemic states, i.e., by representing and sharing strategic knowledge in the overall system.

On the other side, artifacts operational function allows agents to improve the repertoire of actions, i.e., by providing additional means which can be purposively triggered by agents in order to achieve their goals.

Some of the practical benefits of the approach are discussed along with a test-bed scenario –built with CARTAGO– where Jason and Jadex agents are engaged in goal oriented activities relying on the transmission of relevant knowledge and mediated kinds interactions wich are enabled by cognitive artifacts.

Keywords: Goals, Artifacts, BDI, Interaction, Team

Joint work of: Piunti, Michele; Ricci, Alessandro

From Agents to Artifacts Back and Forth: Purposive and Doxastic use of Artifacts in MAS

Michele Piunti (ISTC - CNR - Roma, Italy)

Recent approaches in Multi-Agent Systems are focusing on providing models and methodologies for the design of environments and special purpose tools supposed to scale up complexities. Among others, the Agents and Artifacts (A&A) approach introduced the notion of artifact as first class abstraction providing agents with external facilities, services and coordination medium explicitly conceived for easing their activities. In this paper we analyse A&A systems by focusing on the functional roles played by artifacts. In particular, we here investigate the function of artifacts once they are employed in the context of societies of cognitive agents, i.e. agents capable to reason about their epistemic and motivational states. In this context, a twofold kind of interactions is envisaged. On the one side, artifact representational function allows agent to improve epistemic states, i.e., by representing and sharing strategic knowledge in the overall system (doxastic use). On the other side, artifacts operational function allows agents to improve the repertoire of actions, i.e., by providing additional means which can be purposively triggered by agents to achieve goals (operational use). Some of the outcomes of this approach are discussed along with test cases showing agents engaged in goal-oriented activities relying on the transmission of relevant knowledge and the operations provided by artifacts.

Keywords: MAS, Cognitive Agent, Agents and Artifacts, Interaction, Reasoning, Environments

Joint work of: Piunti, Michele; Ricci, Alessandro

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1642>

Environment Design and Programming in MAS: CARTAGO and (Intelligent Agent) Friends

Alessandro Ricci (University of Bologna, Italy)

In this talk we discuss environment design and programming in Multi-Agent Systems, using CARTAGO as reference programming model and platform. Based on the A&A conceptual model, CARTAGO promotes the design and development of agent environments in terms of open set of artifacts, collected in workspaces, that agents dynamically instantiate, share and use to support their individual and collective work. The talk is about the current state of the research, ongoing work and future directions.

Among the topics planned in the talk:

- the integration of CARTAGO with heterogeneous agent platforms - Jason, JADEX, simpA will be touched in particular - enabling in the overall the development and execution of open and heterogeneous MAS;
- the issue of goal-directed use of artifacts, towards MAS with intelligent agents reasoning about the artifacts to select and use to achieve their goals and/or adapting their computational environment (e.g. by creating new artifacts, linking and manipulating existing ones).

Keywords: Artifacts, CARTAGO, MAS environments, MAS programming, A&A

Joint work of: Ricci, Alessandro; Piunti, Michele

An Organizational Perspective of MAS: Models and Implementations

Jaime Sichman (University Sao Paulo, Brazil)

In this presentation, we will propose to use organizational issues to classify different ways to model and to program multi-agent systems. This classification takes into account mainly (i) the separation between individual agents and the whole system (ii) the level of autonomy of the individual agents with respect to the whole system and (iii) the description level where these issues are taken into account, i.e, design and/or implementation.

Keywords: Agents organizations, autonomy, organizational middleware

Joint work of: Sichman, Jaime; Boissier, Olivier; Hübner, Jomi

A Formal Model of Emotions: Integrating Qualitative and Quantitative Aspects

Bas Steunebrink (Utrecht University, Netherlands)

When constructing a formal model of emotions for intelligent agents, two types of aspects have to be taken into account. First, qualitative aspects pertain to the conditions that elicit emotions. Second, quantitative aspects pertain to the actual experience and intensity of elicited emotions. In this presentation, we show how the qualitative aspects of a well-known psychological model of human emotions can be formalized in an agent specification language and how its quantitative aspects can be integrated into this model. Furthermore, we discuss several unspecified details and implicit assumptions in the psychological model that are explicated by this effort.

Keywords: Emotions, Intelligent Agents

Joint work of: Steunebrink, Bas R.; Dastani, Mehdi; Meyer, John-Jules Ch.

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1644>

Autonomous agent programming: a truth-maintenance system usage

Patrick Taillibert (Thales - Elancourt, France)

Managing uncertainty is an obligation when autonomous agents are concerned; since, because of autonomy, the behavior of the other agents is never completely foreseeable. It is also the case due to the limited knowledge agents have of their environment.

Among the huge set of tools developed so far in computer science for managing uncertainty, truth maintenance systems play an important role. They have been especially successful in diagnosis applications but can be of a more general use in helping problem solver design.

The talk will show how an ATMS can be used in an agent programming language, putting the emphasis on the management of inconsistent beliefs and on the programming of non-synchronous behaviors.

Keywords: ATMS, uncertainty, inconsistency, plan repair

Joint work of: Taillibert, Patrick; Dekoker, Sylvain

Orwell’s Nightmare for Agents? Programming Multi-Agent Organisations

Nick Tinnemeier (Utrecht University, Netherlands)

In this talk I present a programming language that is designed to implement multi-agent organisations. These organisations are developed as separate entities regulating the behaviour of individual agents that interact with the organisation. The focus is on the normative aspect of organisations that are specified in terms of norms being enforced by monitoring, regimenting and sanctioning mechanisms.

Keywords: Programming, Multi-Agent Organisations, Norms, Sanctions, Monitoring

Specifying and verifying commitment-based agent societies using REC+SCIFF

Paolo Torroni (University of Bologna, Italy)

The definition of choreography specification languages for Service Oriented Systems poses important challenges. Mainstream approaches tend to focus on procedural aspects, leading to over-constrained and overspecified models. Because of such a drawback, declarative languages are gaining popularity as a better way to model service choreographies. A similar issue was met in the MAS domain, where declarative approaches based on social semantics have been used to capture the nature of agent interaction without over-constraining their behavior.

This talk presents an integrated framework capable to cover the entire cycle of specification and verification of choreographies, by mixing approaches coming from the Service Oriented Computing and MAS research domains. SCIFF is the underlying logic programming framework for modeling and verifying interaction in open systems. The use of SCIFF brings us two main advantages: (1) it enables capturing within a single framework different aspects of a choreography, ranging from constraints on the flow of messages to effects and commitments resulting from their exchange; (2) it provides an operational model that can be exploited to perform a variety of verification tasks.

Keywords: Multi-agent systems, open systems, interaction protocols, web services, choreographies, declarative specifications, static verification, runtime monitoring and verification, abductive logic programming, sciff, climb, decserflow, commitments, rec, reactive event calculus, computational logic

Goals in Agent Systems: A Unifying Framework

Birna Van Riemsdijk (LMU München, Germany)

In the literature on agent systems, the proactive behavior of agents is often modeled in terms of goals that the agents pursue. We review a number of commonly-used existing goal types and propose a simple and general definition of goal, which unifies these goal types. We then give a formal and generic operationalization of goals by defining an abstract goal architecture, which describes the adoption, pursuit, and dropping of goals in a generic way. This operationalization is used to characterize the discussed goal types.

Keywords: Agent Programming, Goals, Formal Semantics

Full Paper:

<http://www.pst.ifi.lmu.de/%7Eriemsdijk/papers/aamas08.pdf>

See also: M. Birna van Riemsdijk, Mehdi Dastani, Michael Winikoff, Goals in Agent Systems: A Unifying Framework, In Proceedings of the seventh international joint conference on autonomous agents and multiagent systems (AAMAS'08), pp. 713-720. 2008.

Belief update in AgentSpeak-DL

Renata Vieira (PUCRS - Porto Alegre, Brazil)

In previous work (Moreira et al, DALT 2005) we proposed an extension for the belief base of AgentSpeak agents based on Description Logic (DL), aiming at enabling agent oriented programming to cope with recently proposed technologies for the Semantic Web. In such an extension an agent belief base contains the definition of complex concepts, besides specific factual knowledge.

The foreseen advantages are: (i) more expressive queries to the belief base; (ii) a refined notion of belief update, which considers consistency of a belief addition; (iii) flexibility in plan searching allowed by subsumption relation between concepts; and (iv) knowledge sharing in a semantic web context (based on OWL).

Following this proposal an extension of the well know Agent Speak interpreter, Jason, was presented by K lapiscak and Bordini in DALT 2008. Among the interesting open issues is how to deal with the addition of beliefs which violates ontology consistency. In this talk we will discuss this problem related to ABox updating in the context of AgentSpeak-DL.

Joint work of: Moreira, Alvaro; Vieira, Renata

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1641>

Ten Challenges for Normative Multiagent Systems

Leon van der Torre (University of Luxemburg, Luxemburg)

In this paper we discuss the shift from a legal to an interactionist view on normative multiagent systems, examples, and ten new challenges in this more dynamic setting.

Keywords: Normative multiagent systems, normative systems, deontic logic, multiagent systems

Joint work of: Boela, Guido; van der Torre, Leendert; Verhagen, Harko

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1636>