

Formal Approaches to Multi-Agent Systems: Special issue of best papers of FAMAS 2009

This special issue of the Logic Journal of the IGPL includes revised and updated versions of the best work presented at the fourth edition of the workshop Formal Approaches to Multi-Agent Systems, FAMAS'09, which took place in Turin, Italy, from 7 to 11 September, 2009, under the umbrella of the Multi-Agent Logics, Languages, and Organisations Federated Workshops (MALLOW). For an overview of the history and topics of the FAMAS workshop series, we refer the reader to the introduction of the companion special issue of best papers of FAMAS 2007 in this same volume.

Just like its predecessor, research reported in this FAMAS 2009 special issue is very much inspired by practical concerns. This time the authors of all the five selected papers are concerned with knowledge and beliefs in multi-agent settings: How to create a group belief in a fair way from individual plausibility orderings? How to close gaps and resolve ambiguities in a tractable way, when information comes from multiple sources? How to reason about a spatial environment? How to compare the strengths of an agent's beliefs in a principled way? How to decide as efficiently as possible whether a given formula concerning group beliefs is valid? These questions and their answers lead to a multi-faceted and at the same time coherent special issue. We concisely introduce the five articles.

Mechanisms for reaching agreement

In "Protocols for belief merge: Reaching agreement via communication", Baltag and Smets consider an issue often studied in social choice theory: how can a group of agents' individual knowledge and beliefs be merged in a fair way into a single set of group beliefs. The authors advocate a dynamic approach to this problem, based on dynamic epistemic logic with doxastic plausibility orders, that turn out to capture all the agents' individual informational attitudes, from 'soft', potentially fallible, beliefs through conditional beliefs and 'safe' beliefs (in the sense that no future learning of truthful information will force the agent to revise them), up to 'hard' knowledge.

The simplest example of their dynamics is that of distributed 'hard' knowledge, which can be transformed into group knowledge if the agents make sincere public announcements of all their individual knowledge, as advocated by Van Benthem in his 2002 Logic Colloquium paper "One is a lonely number". At the other extreme, for agents' 'soft' beliefs, Baltag and Smets propose to enforce a hierarchy, by giving priority to certain agents' beliefs over others in a so-called radical or lexicographic merge.

Baltag and Smets also consider some more or less democratic intermediate possibilities, such as sincere persuasive public announcements of 'soft' (defeasible) knowledge. They illustrate their definitions with concrete examples. It turns out that an essential role is assigned to the person who sets the agenda and determines the order in which speakers and issues are scheduled.

A comparative approach to beliefs

In the paper “Comparing strengths of beliefs explicitly”, authors Ghosh and De Jongh introduce a novel ordering of formulas: $\varphi \succ_B \psi$ for an agent, if the agent’s strength of belief in φ is greater than that in ψ . Such explicit comparative formulas can be used to express notions such as *plausibility* of φ , meaning that an agent’s strength of belief in φ is greater than that in $\neg\varphi$. Similarly, *disbelief* in φ occurs in situations where an agent’s strength of belief in $\neg\varphi$ is greater than that in φ .

Ghosh and De Jongh show how this added expressivity to the standard logic of beliefs helps to model common decision situations. The authors investigate both the resulting comparative belief ordering and a different but related plausibility ordering reminiscent of David Lewis’ sphere systems for counterfactuals. Then they provide a sound and complete axiomatization for the single-agent case. The four different logics introduced in this paper to express different notions of belief and plausibility are then subjected to a conclusive comparative expressiveness study. The authors also compare their proposal to the notion of ‘safe belief’ discussed by Baltag and Smets in this special issue. Finally, the authors extend their approach to multi-agent situations.

Taking geometry seriously

Balbani, Gasquet and Schwarzentruher aim to develop a multi-agent epistemic logic based on spatial geometric semantics in their contribution “Agents that look at one another”. Interestingly, the paper was originally inspired by didactical considerations: In a given concrete situation in which agents and lamps are placed on a line, what do agents know about lamps and about the knowledge of the other agents of these? Such concrete reasoning based on what agents can and cannot see, may help students to learn abstract aspects of epistemic logics, such as higher-order knowledge and the effects of public announcements.

In the current journal paper, the authors take their considerations much further than this ‘Lineland’. They argue that multi-agent logics often have a very abstract semantics, in seeming contrast with the fact that real agents like robots in rescue situations or virtual characters in video games are embedded in their spatial environment and, importantly, can see only certain portions of their surroundings. The authors introduce multi-agent logics for which they define the semantics straightforwardly from geometrical notions. In their set-up, possible worlds are given by means of the positions in \mathbb{R}^n occupied by agents and the sections of \mathbb{R}^n that they can see. Accessible states for an agent are then naturally construed as those states that the agent can imagine to be compatible with what it can currently see. The authors investigate several formal properties of the systems they introduce for different dimensions, such as expressivity, axiomatization and completeness, as well as the complexity of model checking problems.

Decision procedures for multi-agent systems

Ajspur, Goranko and Shkatov, in their contribution “Tableau-based decision procedure for the multiagent epistemic logic with all coalitional operators for common and distributed knowledge”, take a second look at the coalitional multi-agent epistemic logic. This logic extends standard epistemic logic by adding modalities expressing distributed and common knowledge among all possible teams of agents in the language.

The authors provide an elegant tableau-based decision procedure for this logic, in which tableaux are built up incrementally. They prove that the procedure is sound and complete, and show that the decision procedure works in exponential time, as expected for epistemic logics with common knowledge. They show that the procedure is more efficient than the usual one based on “maximal” tableaux, as presented in Halpern and Moses’ influential 1992 paper “A guide to completeness and complexity for modal logics of knowledge and belief”. Nevertheless, Ajspur, Goranko and Shkatov suggest several ways in which their decision procedure can be made even more efficient, for example by constructing cut-free tableaux. They also suggest several lines of extension, most importantly to temporal versions of the logics on linear and branching time, as well as to Alternating-time temporal epistemic logic (ATEL).

Taking uncertainty on board without losing efficiency

The article “How an agent might think” by Andrzej Szalas steps away from the modal logics investigated in the other articles in this special issue, and for good reason. The author argues that agents get their information from multiple sources, which may lead to incomplete but also to inconsistent information. Usually such issues are solved by applying non-monotonic or multi-modal reasoning techniques to fill the gaps and resolve ambiguities. These logics, however, lead to decision problems that are far from tractable, therefore not applicable in real-world intelligent systems.

Instead, the author proposes a natural four-valued paraconsistent semantics, and explains in the process why the different four-valued lattice introduced by Belnap in his 1977 paper “How a computer should think” leads to unwanted consequences when used in practical applications. This paper provides an extension of the rule query language 4QL, which Szalas introduced with Maluszyński and extended to a multi-agent context with Dunin-Kępicz in recent years. In the current extension, a number of constraints of 4QL are relaxed, for example, in the new system 4QL⁺, multisource queries to heterogeneous databases can be posed. It is shown that these extensions, while allowing wider applicability in agent systems, do not endanger the nice tractability features of 4QL: 4QL⁺ still has PTIME data complexity, and manages to capture PTIME queries on ordered structures.

Acknowledgments

We would like to thank all the people who helped to bring about FAMAS’09. First of all, we thank all speakers for ensuring a diverse and interesting workshop. The

FAMAS website would not exist without the help of Michał Ślizak in Warsaw. Special thanks are due to the members of the program committee for their professionalism and their dedication to select papers of quality and to provide authors with useful, constructive feedback during the in-depth reviewing process: Hans van Ditmarsch, Andreas Herzig, Wiebe van der Hoek, Wojtek Jamroga, John-Jules Meyer, Wojciech Penczek, David Robertson, Jeff Rosenschein, Carles Sierra, Andrzej Szalas, Onn Shehory. We would also like to express our gratitude to the anonymous reviewers of this special issue who have written constructive reviews with many useful detailed suggestions for the authors, often on several versions.

Warszawa
Groningen

Barbara Dunin-Kępicz
Rineke Verbrugge

December / January, 2012