The Knowledge Engineering Review, Vol. 17:3, 295–301. © 2002, Cambridge University Press
DOI: 10.1017/S0269888902000504 Printed in the United Kingdom

Practical and theoretical innovations in multi-agent systems research

MARK D'INVERNO,[1] MICHAEL LUCK[2] and UKMAS 2001 Contributors*

[1] *Cavendish School of Computer Science, University of Westminster, UK*
[2] *Department of Electronics and Computer Science, University of Southampton, UK*

1 Introduction

UKMAS has now been running for six years, in 1996 and 1997 under the heading of FoMAS (Foundations of Multi-Agent Systems) both organised by Michael Luck at Warwick University and then subsequently in its current incarnation, UKMAS, first by Michael Fisher at Manchester Metropolitan University then by Chris Preist at Hewlett Packard Laboratories, Bristol and finally by Mark d'Inverno at St Catherine's College, Oxford in 2000. After the success of the workshop last year at St Catherine's in providing an excellent opportunity for academics and industrialists to come together to discuss current work and directions in the multi-agent systems field, it was decided by the steering committee to use St Catherine's once again as the venue for UKMAS 2001. The workshop was sponsored by the Engineering and Physical Sciences Research Council and by AgentLink, the European Commission's IST-funded Network of Excellence for Agent-Based Computing.

The UKMAS workshops are concerned both with the dissemination of recent research within multi-agent systems and with the provision of an appropriate forum for debate and discussion between students, academics and industrialists alike. As with UKMAS 2000, this year's event attracted around 80 participants and the format of this meeting was typical of previous years (Luck, 1997; Doran *et al*., 1997; d'Inverno *et al*., 1997; Fisher *et al*., 1997; Luck *et al*., 1998; Aylett *et al*., 1998; Binmore *et al*., 1998; Aylett *et al*., 2000; Beer *et al*., 1999; Decker *et al*., 1999; Chattoe *et al*., 2000; Rana *et al*., 2000; Alonso *et al*. 2001; d'Inverno *et al*., 2001). There were two excellent invited presentations, the first by Samson Abramsky, the Christopher Strachey Professor of Computing at the Oxford University Computing Laboratory, who gave a talk entitled "From computation to interaction", and the second from Sarit Kraus of Bar Ilan University, with a presentation on "Real-time cooperation and negotiation". In addition, there were 18 paper presentations and a panel session entitled "Agents in e-commerce", which was chaired by Professor Nick Jennings from Southampton University.

2 The first day

Samson Abramsky of Oxford University started the first day with a talk entitled 'From computation to interaction'. As information technology has progressed from batch processing through multi-tasking operating systems to distributed systems and on to today's Internet, the focus has correspondingly

* UKMAS 2001 Contributors: Samson Abramsky; Wamberto Vasconcelos, Jordi Sabater, Carles Sierra and Joaquim Querol; Pietro Panzarasa and Nick Jennings; Peter McBurney and Simon Parsons; Colm O'Riordan; Rogier van Eijk, Frank de Boer, Wiebe van der Hoek and John Jules Ch Meyer; Maria Fasli; Alessio Lomuscio and Marek Sergot; Martijn Schut and Mike Wooldridge; Tim Norman and Chris Reed; Nick Jennings; Martin Kollingbaum and Tim Norman; Andrew Byde; Jean-Luc Koning and Marc-Philippe Huget; Neil Johnson, David Lamper, Paul Jefferies, Michael Hart and Sam Howison; Sarit Kraus; Anthony Karageorgos, Simon Thompson and Nikolay Mehandjiev; Rodrigo Machado and Rafael Bordini; Christoph Oechslein, Franziska Klugl and Frank Puppe; Ronald Ashri and Michael Luck; Hong Zhu; Daniel Kudenko, Mathias Bauer and Dietmar Dengler; Michael Schroeder and Penny Noy; Virupatchan Maheswaran and Nick Antonopoulos; John Bigham, Damien Ryan, Wai Sum Thong and Laurie Cuthbert.

shifted, from stand-alone *programs* to *systems*, and from *computation* to *interaction*. The "first-generation" models of computation gave an account of stand-alone programs computing functions from inputs to outputs. These are now seen as (very) special cases of a wider class of behaviours, in which the components of a complex system of concurrently executing agents interact to achieve some global effect. On a "macro" scale, this description evidently fits today's distributed transactions across the Internet, but on a "micro" scale it is equally true of how functional computations are ultimately realised, whether in software or in hardware. Ultimately, it seems that all computation can be resolved into the interactions of large numbers of very simple agents; the complex behaviour of the overall system is an "emergent property" of these interactions.

Samson stated that the challenge has then been to find models of computation which take interaction as their basic ingredient. These models should combine sufficient expressive power to yield faithful descriptions of the phenomena arising from contemporary developments in IT, with sufficient mathematical structure and tractability to provide a basis for the formal analysis of such systems. The fundamental character of this challenge is indicated by the following observation. The notions of *function*, *set* and *algorithm* were available "off-the-shelf" from mathematics and logic for use in computer science. By contrast, there is no adequate pre-existing theory of *processes*, *interaction*, *information flow* and so on, on which "second-generation" models can build. Rather, the most basic task of such models is precisely to give a non-trivial analysis of these concepts.

Samson subsequently outlined important developments from linear logic, geometry of interaction and game semantics, which have made a distinctive contribution to these issues. In particular, Game Semantics, in which a system is seen as interacting with its environment by playing a game with it, has proved to be a powerful and versatile tool. While somewhat akin to the way games are used to model interactions between rational agents in economics, game semantics pays much more attention to the fine structure of the interactions themselves, rather than focusing mainly on the outcomes of the plays. While game semantics draws on previous work on games in logic, it pays much more attention to the mathematical structure of categories of games, and to issues of *compositionality* and *syntax-independence* – which are key to the analysis of interaction. This issues were expertly explored in a lively and illuminating talk.

In the first of four paper sessions Wamberto Vasconcelos from the University of Edinburgh described joint work with Jordi Sabater, Carles Sierra and Joaquim Querol (all from the Artificial Intelligence Research Institute (IIIA) in Spain) focused on developing agent systems for electronic institutions using skeletal designs and semi-automatic agent development. The scenario used is one in which agents are designed to follow an electronic institution using a formalism to specify open agent organisations. In this approach, an initial design pattern is automatically extracted from a given electronic institution, which is then offered to programmers willing to develop agents to perform effectively in electronic institutions. A key advantage of this approach is that it supports developers when modifying the initial simple design pattern into more sophisticated programs. Next, Pietro Panzarasa, in collaboration with Nick Jennings (both of the University of Southampton), proposed a need for a new foundation of the science of multi-agent systems. In this work they argue that such a science should be grounded, theoretically, on a richer conception of sociality and, methodologically, on the extensive use of computational modelling for real-world applications and social simulations. In the presentation, a model of multi-agent systems was proposed that reflects a fully explicated conception of cognition both at the individual and the collective level. Subsequently, the mechanisms and principles underpinning the model were detailed with particular emphasis on the contributions provided by contemporary organisation theory.

Peter McBurney then presented collaborative work with Simon Parsons (both at the University of Liverpool) concerning a new semantics for dialogue game protocols used in autonomous agent interaction. Dialogue games have been studied by philosophers since the time of Aristotle and have recently found application in artificial intelligence as the basis for protocols for interactions between autonomous software agents. For instance, game protocols have been proposed for agent dialogues involving team formation, persuasion, negotiation and deliberation. There is as yet, however, no formal mathematical theory of dialogue game protocols with which to compare two protocols or to study their

formal properties. As a step towards such a theory, a new geometric semantics for such protocols has been developed in this work and, with it, a notion of equivalence between two protocols defined. This algebraic property of protocol equivalence was demonstrated and it was argued that this result has implications for the design and evaluation of agent dialogue-game protocols in general.

Next, Colm O'Riordan from the National University of Ireland, Galway, presented and discussed some results in game-theoretic modelling of agent-based systems. He first gave a brief discussion of cooperation in multi-agent systems and then concentrated on the prisoner's dilemma and the iterated prisoner's dilemma as an abstraction of some of the salient features involved in repeated interactions of agents in a multi-agent system. Subsequently, he presented the results of a new family of strategies in the setting of the iterated prisoner's dilemma, and discussed the general applicability to multi-agent systems.

The final presentation in the session was from Rogier van Eijk concerning a verification framework for agent communication that had been developed with Frank de Boer, Wiebe van der Hoek and John-Jules Meyer all at Utrecht University in the Netherlands. The verification method for the correctness of multi-agent systems was in relation to an existing communication framework called ACPL (Agent Communication Programming Language) presented at previous UKMAS workshops. The computational model of ACPL consists of an integration of the two different paradigms of CCP (Concurrent Constraint Programming) and CSP (Communicating Sequential Processes). The constraint programming techniques are used to represent and process information, whereas the communication mechanism of ACPL is described in terms of the synchronous handshaking mechanism of CSP. A verification method for ACPL defined in terms of an integration of the verification methods for CCP and CSP has been developed. In addition it has been possible to prove formally the soundness of the method and some issues concerning its completeness were also discussed.

After lunch, Maria Fasli of the University of Essex began the afternoon's paper session by presenting a formalisation of commitments, roles and obligations in multi BDI-agent systems. A formal analysis of general obligations and relativised-to-one obligations from a bearer to a single counterparty was described and obligations were examined in the context of strong realism, realism and weak realism for BDI agents. Fasli subsequently described how relativised-to-one obligations could arise as a result of social commitments and the adoption of roles. In this framework, if an agent adopts a role, then the role is associated with one or more social commitments. Social commitments give rise to relativised obligations and, consequently, roles, social commitments and relativised obligations are interwoven.

In the next presentation Alessio Lomuscio described work with Marek Sergot (both at Imperial College, London) on an extension of the formalism of interpreted systems by Halpern and colleagues to model correct behaviour of agents. This semantic model allows for the representation and reasoning about states of correct and incorrect functioning behaviour of the agents, and of the system as a whole. Lomuscio described how the axiomatisation of this semantic class could take effect by mapping it into a suitable class of Kripke models. The resulting logic, KD45, is a stronger version of KD, the system often referred to as Standard Deontic Logic. The logic was explicated using some simple examples.

In the third presentation of the afternoon, Martijn Schut gave a talk on collaborative work with Mike Wooldridge (both at the University of Liverpool). He presented a framework that enables a Belief–Desire–Intention (BDI) agent to dynamically choose its intention reconsideration policy in order to perform optimally in accordance with the current state of the environment. This framework integrates an abstract BDI agent architecture with a decision-theoretic model for discrete deliberation scheduling. As intention reconsideration determines an agent's commitments to its plans, this work increases the level of autonomy in agents as it pushes the choice of commitment level from design time to run time. This makes it possible for an agent to operate effectively in dynamic and open environments whose behaviour is not known at design time. Following a precise formal definition of the framework, an empirical analysis that evaluates the run-time policy in comparison with design-time policies was detailed.

Tim Norman (from the University of Aberdeen) then gave a talk on a new theory of group delegation and responsibility developed with Chris Reed (from the University of Dundee), two key topics in the

theory multi-agent systems. The act of delegating a task by one agent to another can be carried out by the performance of one or more communicative acts. Such acts may not only be directed to another individual but to a group of agents. In this presentation, the semantics of imperatives are explored with reference to extant logics of agentative action, and in the context of the referent of an imperative being either an action or a state of affairs. The particular case of issuing an imperative to a group of individuals was described in detail.

The first day ended with Nick Jennings from the University of Southampton chairing a panel session on the use of agents in e-commerce. The panelists were Andrew Byde from Hewlett Packard Labs, Bristol, Marc-Philippe Huget from the University of Liverpool, David Lamper from Oxford University and Martin Kollingbaum from the University of Aberdeen. All panelists were asked three broad questions: their favourite current example of an agent system in e-commerce (which clearly illustrates the value-added nature of the agentification), the most compelling application of or role for agents in future-generation e-commerce systems and the key impediments (technical and/or social) to the widespread adoption of agent-mediated electronic commerce.

The panelists were invited to show how their work related to the issues identified in the last question. Martin Kollingbaum proposed a mechanism – supervised interaction involving three roles (addressee, counter-party and authority) – designed to provide the web of trust necessary for successful deployment of agent-mediated electronic markets. Andrew Byde presented algorithms for agents participating in multiple simultaneous auctions for a single private-value good including stochastic dynamic programming to derive formal methods for optimal algorithm specification. Then Marc-Philippe Huget presented a full development life cycle for the engineering of an electronic commerce interaction protocol. Finally, David Lamper reported on a technique based on multi-agent games which has potential use in the prediction of future movements of financial time-series.

3 The second day

Sarit Kraus gave a talk entitled "Real time negotiation and cooperation" that covered two projects, both concerned with the development of agents that can cooperate under time pressure. Whilst the first project concerned cooperative agents, the second related to self-interested agents that could benefit from cooperation. First, the Distributed Dispatcher Manager (DDM), a system for managing large collections of dynamically changing tasks, was described. In this work, tasks are distributed over large geographic areas and teams consist of very large groups of mobile and cooperative agents which have direct access to only local information about their immediate environment. The DDM's contributions include real-time processes for combining partial results to form an accurate global solution, increased system fault tolerance and scalability to very large task and agent problem domains. The second project described the development of an automated agent that can negotiate efficiently with humans. In this work, the environment is characterised by two negotiators, time constraints, deadlines, full information and the possibility of opting out. The agent can play either role, with communications between agent and human taking place using a semi-formal language. The model used in constructing the agent is based on a formal analysis of the scenario using game-theoretic methods and heuristics for argumentation. The agent receives messages sent by humans, analyses them and responds. Experimental results were presented of simulations of a fishing dispute between Canada and Spain which indicated that the agent played at least as well as, and in the case of Spain, significantly better, than a human player.

After coffee, Simon Thompson from the Intelligent Agents Research Group, BT, discussed some joint work with Anthony Karageorgos and Nikolay Mehandjiev from UMIST, Manchester. The position of this group is that designing an agent organisation is a complex process involving defining the structural relationships among agents, the lines of inter-agent communication and the agent functionality. Existing approaches to agent organisation design are difficult to apply in practice since they require designers to make decisions while working at a low level of abstraction. This work is concerned with designing agent organisations in a practical and effective manner by proposing to semi-automate the organisational design process. This semi-automatic approach enables agent-system

designers to reason at a high abstraction level and conveniently reuse previous design decisions. It used role modelling and a role algebra to capture a number of basic relations among roles. The role algebra's semantics are formally defined using a two-sorted algebra and the role's use demonstrated by an example drawn from a case study involving telephone repair service teams.

Next, Rafael Bordini discussed an ongoing project with Rodrigo Machado (both at the Universidade Federal do Rio Grande do Sol in Brazil) on running agents built using the AgentSpeak(L) language on top of an agent toolkit called SIM_AGENT. This work is arguably the first successful attempt at running programs in AgentSpeak(L), which is a programming language for BDI agents, created by Rao. The authors argue that AgentSpeak(L) is an elegant and neat notation for a BDI programming language and could establish a turning point in the practice of implementing cognitive multi-agent systems. However, because there is no associated interpreter or compiler, AgentSpeak(L) has been neglected, as have other agent-oriented programming languages, by the multi-agent systems community. This paper shows a way of turning AgentSpeak(L) agents into running programs with Sloman's SIM_AGENT toolkit. It also outlines the effectiveness of AgentSpeak(L) as an agent programming language as well as showing the power of the SIM_AGENT toolkit.

In the third of the morning's paper presentations Christoph Oechslein presented work developed with Franziska Klugl and Frank Puppe at the University of Würzburg. He argued that whilst developing multi-agent simulations appeared to be rather straightforward (as active entities in the original correspond to active agents in the model and so plausible behaviours can be produced relatively easily), for real-world applications agents must satisfy some requirements concerning verification, validation and reproducibility. Using a standard framework for designing a multi-agent model one can gain further advantages like fast learnability, wide understandability and possible transfer. In this presentation, Oechslein showed how UML can be used to specify behaviour-oriented multi-agent models. The focus here is on activity graphs and the representation of different forms of interaction in these graphs. He argued that UML provides a means for specifying model invariants and constraints on static and dynamic relations between agents in order to support quality assurance.

Next, Ronald Ashri described an ongoing project developing the Paradigma system, based on the Smart agent framework, with Michael Luck at the University of Southampton. It is clear by now that the take-up of agent technologies and the wide use of such technologies in open environments depends on the provision of appropriate infrastructure to support the rapid development of applications. In his presentation Ashri argued that elements required for the development of infrastructure span three different fields which, nevertheless, have a great degree of overlap. Middleware technologies, mobile agent and intelligent agent research all have significant contributions to make towards a holistic approach to infrastructure development, but it is necessary to make clear distinctions between the requirements at each level and explain how they can be integrated so as to provide a clearer focus and allow the use of existing technologies. This view of the requirements for infrastructure to support agent-based systems has been formed through experience with developing an agent implementation environment based on a formal agent framework. In order to provide support to developers, this infrastructure must address conceptual concerns relating to the different types of entities and relationships between agent and non-agent entities in the environment.

Then, Hong Zhu from Oxford Brookes University described his ideas on formally specifying multi-agent systems. He argued that one of the most appealing features of multi-agent technology is its natural way of modularising a complex system in terms of multiple, interacting and autonomous components. As a natural extension of classes, he proposed *castes* (which represent a set of agents of common structural behavioural characteristics) and introduced them into an existing formal specification language to provide a language facility that provides modularity. He showed how the caste facility could be employed to specify multi-agent systems so that the notion of roles, organisational structures of agent societies, communication, collaboration protocols and so on can be naturally represented. Daniel Kudenko from the University of York then discussed ongoing work with Mathias Bauer and Dietmar Dengler from the German Research Center for Artificial Intelligence concerning online purchasing using user-modelling and multi-agent systems negotiation and argumentation. To date, online catalogue purchase assistants have mainly been investigated from a

single-agent perspective, where only the interests of a single user are taken into account. In this work, this scenario is extended to look at cases where more than one user is involved in the purchasing, and therefore many (potentially conflicting) interests have to be considered. An overview of a system that assists a group of users to reach a joint decision on an online catalogue purchase was described. This is achieved by acquiring individual user models and using them to simulate negotiations that are subsequently presented to the users for further explanation and argumentation.

In the first presentation of the fourth and last paper session Michael Schroeder described work with Penny Noy at City University concerned with capturing features of complex agent systems using multivariate data. The focus of their work is on methods that reduce the dimensions of this data through matrix transformations and then visualise the entities in the lower dimensional space. An approach describing agent similarities through distances, which are then visualised by multi-dimensional scaling techniques, was analysed and its shortcomings discussed. An alternative, which applies principal component analysis and subsequent visualisation directly to the data, and has been implemented in the Space Explorer tool, was then detailed. Schroeder identified four categories of data, which capture interaction, profiles time series, and combinations of these three. Finally, it was shown how to employ them for various agent types such as communicating, mobile, personal, interface, information and collaborating agents.

Next, Virupatchan Maheswaran described work with Nick Antonopoulos at the University of Surrey on using agents to support interoperability between heterogeneous GRIDs. New advanced network-based applications are emerging that require services not provided by today's Internet and Web environments but instead need a GRID that both integrates new types of resource into the network and provides enhanced "middleware" services. The initial step taken in this work is to outline and study the existing GRID frameworks, focusing on their main purpose and the extensibility of the components. Subsequently, a peer-peer structure is proposed, which will enable GRID interoperation in the line of information search and retrieval.

Finally, John Bigham described the use of agents in managing third-generation mobile networks, a project undertaken with Damien Ryan, Wai Sum Thong and Laurie Cuthbert at Queen Mary, London. Specifically, the work is concerned with the use of intelligent agents to provide an efficient, scalable and robust mechanism to manage 3G networks in a deregulated marketplace containing several service providers, network operators and customers. In 3G mobile networks, management of the radio resource is more complex than in existing 2G networks; higher bandwidth services requiring better quality of service, liberalisation of the marketplace and technical issues with the radio access method all contribute to this problem. In this presentation, details of the reactive components in two of the agents were provided.

4 Summary

Once again UKMAS 2001 provided an ideal opportunity for agent researchers, practitioners and students to come together in an informal and lively but structured environment in order to present, debate and discuss current issues in the development of multi-agent research and development within the UK. For further details of the workshop and future events, visit `http://www.ukmas.org`.

References

Aylett, R, Brazier, F, Jennings, N, Luck, M, Preist, C and Nwana, H, 1998, "Agent systems and applications" *Knowledge Engineering Review* **13**(3) 303–308.

Aylett, R, Dautenhahn, K, Doran, J, Luck, M, Moss, S and Tennenholtz, M, 2000, "Can models of agents be transferred between different areas?" *Knowledge Engineering Review* **15**(2) 197–203.

Alonso, E, d'Inverno, M, Kudenko, D Luck, M and Noble, J, 2001, "Learning in multi-agent systems" *Knowledge Engineering Review* **16**(3) 277–284.

d'Inverno, M and Luck, M, 2001, "Multi-agent systems research into the 21st century" *Knowledge Engineering Review*, **16**(3) 271–275.

Beer, M, d'Inverno, M, Jennings, N, Luck, M, Preist, C and Schroeder, M, 1999, "Negotiation in multi-agent systems" *Knowledge Engineering Review* **14**(3) 285–289.

Binmore, K, Castelfranchi, C, Doran, J and Wooldridge, M, 1998, "Rationality in multi-agent systems" *Knowledge Engineering Review* **13**(3) 309–314.

Chattoe, E, Dautenhahn, K, Dickinson, I, Doran, J and Vulkan, N, 2000, "A multidisciplinary perspective on multi-agent systems" *Knowledge Engineering Review* **15**(3) 285–292.

Decker, K, Fisher, M, Luck, M, Tennenholtz, M and UKMAS'98 Contributors, 1999, "Continuing research in multi-agent systems" *Knowledge Engineering Review* **14**(3) 279–283.

D'Inverno, M, Fisher, M, Lomuscio, A, Luck, M, de Rijke M, Ryan, M and Wooldridge, M, 1997, "Formalisms for multi-agent systems" *Knowledge Engineering Review* **12**(3) 315–321.

Doran, JE, Franklin, S, Jennings, NR and Norman, TJ, 1997, "On cooperation in multi-agent systems" *Knowledge Engineering Review* **12**(3) 309–314.

Luck, M, 1997, "Foundations of multi-agent systems: issues and directions" *Knowledge Engineering Review* **12**(3) 307–308.

Luck, M, d'Inverno, M, Fisher and FoMAS'97 contributors, 1998, "Foundations of multi-agent systems: techniques, tools and theory" *Knowledge Engineering Review* **13**(3) 297–302.

Rana, O, Preist, C and Luck, M, 2000, "Progress in multi-agent systems research" *Knowledge Engineering Review* **15**(3) 293–301.