*The Knowledge Engineering Review*, Vol. 16:3, 271–275. © 2001, Cambridge University Press
DOI: 10.1017/S0269888901000169 Printed in the United Kingdom

# Multi-agent systems research into the 21st century

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

There is little doubt that the strength and breadth of UK research into multi-agent systems continues to grow as we move into the new millennium. In the middle of an extremely cold December in 2000, the Third UK Workshop on Multi-Agent Systems (UKMAS 2001) was held at St Catherine's College, Oxford. This was the fifth such meeting in as many years, generously sponsored by EPSRC, FIPA (The Foundation for Intelligent Physical Agents) and Hewlett Packard.

UKMAS 2000 saw over 80 participants at the workshop, which has two broad aims: to facilitate the dissemination of recent research within the multi-agent system community from both academia and industry and to encourage debate in this wide-ranging area. 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), with invited presentations from Craig Boutilier of the University of Toronto, Canada, and Carles Sierra of the Artificial Intelligence Research Institute in Spain, several paper presentations and two panel sessions. The first panel was concerned with multi-agent-systems learning, and the second with a retrospective assessment of what has been achieved in the last five years, acknowledging the length of time since the first workshop in the series.

## 2 The first day

The first day of the workshop began with an invited talk from Craig Boutilier of the University of Toronto, Canada entitled "Sequential optimality and coordination in multiagent systems". Clearly, coordination of agent activities is a key problem in multi-agent systems but, set in a larger decision-theoretic context, the existence of coordination problems leads to difficulty in evaluating the utility of a situation. This in turn makes optimal policies for sequential decision processes problematic. Boutilier proposed a method for solving sequential multi-agent decision problems by allowing agents to reason explicitly about specific coordination mechanisms.

In his talk, he described an extension of value iteration in which the system's state space is augmented with the state of the coordination mechanism adopted, allowing agents to reason about the short- and long-term prospects for coordination and the long-term consequences of (mis)coordination, and to make decisions to engage or avoid coordination problems based on expected value. The value of Boutilier's model can be demonstrated using well-known learning mechanisms such as fictitious

* UKMAS 2001 Contriubutors: C Boutilier; C Sierra; S Chong and K Liu; A Bazzan, R Bordini and G Andrioti, R Vicari and J Wahle; M Beer, I Anderson and and W Huang; M Fisher and C Ghidini; M Schut and M Wooldridge; S Fatima; C Dixon, M Fisher and A Bolotov; W de Vries, W van der Hoek, J-J Meyer and F de Boer; T Norman and C Reed; C Preist, C Bartolini and I Phillips; S Poslad, P Buckle and R Hadingham; J Bigham, L Tokarchuk and L Cuthbert; R Ashri and and M Luck; M Wooldridge and A Lomuscio, K Clark and P Robinson; J Doran; O Rana and L Moreau; U Hustadt, C Dixon, R Schmidt, and M Fisher, J-J Meyer and W van der Hoek.

https://doi.org/10.1017/S0269888901000169 Published online by Cambridge University Press

play as an example coordination protocol. As well as illustrating the benefits of mechanism generalisation Boutilier went on to describe some directions in which his work could be extended, including as a basis for the design of coordination mechanisms and learning protocols.

In the first of the paper sessions, Samuel Chong from Staffordshire University described collaborative work with Kecheng Liu developing social software agents for e-commerce systems. He argued that if software agents are autonomous, then it is critical to consider social issues when modelling agent-based e-commerce systems in order to obtain a functional system that truly supports business activities. For this purpose, the DEON methodology was developed based on semiotic principles, and the talk outlined how DEON contributes to the understanding of modelling social obligations of human agents as well as software agents. The methodology also focuses on the descriptive qualities of semiotics that can be useful during the design of interfaces for software agents. The talk undertook to demonstrate how the DEON methodology can offer a unifying framework for identifying the main building blocks of software agents.

Next, Rafael Bordini described a collaborative project (between Universidade Federal do Rio Grande do Sol in Brazil and the Physics of Transport and Traffic at Gerhard-Mercator University in Germany) with Ana Bazzan, Gustavo Andrioti, Rosa Vicari and Joachim Wahle that concerned how, in our day-to-day behaviour, we often have to face binary decisions where we seek to take the minority's choice as, for example, in traffic scenarios where we have to choose between similar alternative routes. Bordini described recent agent coordination mechanisms in a binary decision model known as the Minority Game. By extending this model, different personalities that model certain types of human behaviour have been considered, and simulations of different populations, which contain these personalities, have been built. It was found that there is one personality that performs better than the average: the wayward personality. Bordini argued that this work provides an insight into the impact of commuter behaviours as well as addressing related issues in traditional traffic simulation.

In the last of the morning sessions, Martin Beer described a project at the University of Liverpool which uses an agent system to demonstrate the practicality of the INCA (Intelligent Community Support for the Elderly) architecture. The INCA architecture is intended to integrate a number of autonomous systems – home monitoring, community alarms, and care management systems – using agent technology to build effective coordinated care systems. Since these systems not only contain information relevant to community care, but also all the other activities that one would be unwilling to make available to other parties, the actual management of community care has remained primarily outside the role of current systems. Beer's system was built using the ZEUS agent-building toolkit as the basis for the development of a "benchtop" demonstrator to show that the INCA architecture is both scalable to realistic activity levels, and integrates fully and effectively with existing computer systems in the various agencies involved, without loss of autonomy and security.

Michael Fisher began the afternoon session by speaking about joint work with Chiara Ghidini (both were then of Manchester Metropolitan University) concerning resource-bounded agents and, more specifically, the amount of reasoning they are able to carry out in different situations. The bounds an agent faces are rarely static, and often depend on the situation in which the agent finds itself. For example, at a certain time in a football game, a player might have plenty of time to "think" and decide what to do next. However, if the game is going to end imminently and a quick decision is required, the same player may have relatively little time to "think". In his talk, Fisher described a logic framework for both the logical specification and execution of agents in which dynamic agents that have resource bounds varying over time can be represented.

Another related problem in the design of deliberative agents such as Belief-Desire-Intention (BDI) agents, is that of finding an appropriate policy for intention reconsideration: to deliberate only when necessary. Martijn Schut discussed collaborative work with Michael Wooldridge at Liverpool University concerning the notion of commitment in complex environments. Previously, Kinny and Georgeff had investigated the effectiveness of several such reconsideration policies, and demonstrated that, in general, there is no one best approach, and different environments demand different intention reconsideration strategies. Schut described the relationship between the effectiveness of an agent and its intention reconsideration policy in different environments, and presented an empirical evaluation of

https://doi.org/10.1017/S0269888901000169 Published online by Cambridge University Press

the performance of different reconsideration strategies in environments that are to varying degrees dynamic, inaccessible and non-deterministic.

Also from Liverpool University, Shaheen Fatima proposed an adaptive organisational policy known as TRACE (Task and Resource Allocation in a Computational Economy), incorporating task and resource allocation for multi-agent systems that operate under time constraints and load variations. General multi-agent systems comprise several problem-solving organisations and any task allocation protocol takes requests and plans, and allocates subtasks to agents within an organisation. As requests arrive arbitrarily, at any instant, some organisations could have surplus resources while others could become overloaded. In order to minimize the number of lost requests caused by an overload, the allocation of resources to organisations is changed dynamically by the price-directed resource allocation protocol. Simulation results show that TRACE exhibits high performance despite unanticipated changes in the environment.

Next, collaborative work between Clare Dixon, Michael Fisher and Alexander Bolotov (then all at Manchester Metropolitan University) concerning a resolution-based proof system for a Temporal Logic of Possible Belief was presented. This logic represents a combination of the branching-time temporal logic CTL and the modal logic KD45. Since such combinations of non-classical logics are often used in agent theories for specifying complex properties of rational agents, the resolution system presented here provides an important basis for the verification of such specifications.

In the last talk of the paper session, Wieke de Vries described work at Utrecht University with Wiebe van der Hoek, John-Jules Meyer and Frank de Boer that proposes a new operational model for agents, which includes some aspects of agent systems often neglected. The proposed model formalises the intuitive notion of multiple agents interacting with their environment, through cycles of sensing, reasoning and acting, and has several flexible features. As a result, it can be used to analyse the behaviour of a diversity of agent systems. The first day ended with a panel discussion on agent learning, details of which can be found in the paper in this issue.

## 3 The second day

Carles Sierra, of the Artificial Intelligence Research Institute in Spain, opened the second day with a lively and interesting invited presentation entitled “Agent-mediated interaction: from auctions to negotiation and argumentation”. He began his talk by describing a number of techniques that have been used to structure the relation between customers and sellers over the Internet. Users have been allowed to specify the attributes of goods to be acquired, and their behaviour has been modelled through profiles. However, most developments up to now are centred on a passive web query type of interaction. Sierra surveyed richer and more flexible ways of interaction by means of *mediators*, from the most simple auction mechanisms to argumentative protocols passing through different negotiation techniques. He summarised by stating that because of the massive expansion of all sorts of mobile devices, a deregulation of power and telecom markets, a progressive distribution of computer functionalities and an increase in context perception by computers, secure and well-understood agent interactions will become essential.

Chris Reed of Dundee University gave the first talk in the morning’s paper session, on collaborative work with Tim Norman of Aberdeen University on delegation and responsibility in multi-agent systems. In multi-agent systems, agents may decide to delegate tasks to others. The act of delegating a task by one autonomous agent to another can be carried out by the performance of one or more imperative communication acts. The semantics of imperatives are specified using a language of actions and states. Then, Chris showed how their model can be used to distinguish between the whole-hearted and mere extensional satisfaction of an imperative, and how this may be used to specify the semantics of imperatives in agent communication languages.

In a more applied context, the next speaker was Chris Preist, who discussed an ongoing project at Hewlett-Packard Laboratories with Claudio Bartolini and Ivan Phillips. This is concerned with the design of algorithms for agents to use when participating in multiple simultaneous English auctions, when aiming to purchase multiple goods. Preist first presented a coordination algorithm, which ensures

https://doi.org/10.1017/S0269888901000169 Published online by Cambridge University Press

the agent places appropriate bids in the different auctions to buy exactly the right number of goods. This algorithm is then combined with another algorithm to determine what maximum bid an agent should place in an auction that is about to terminate. Here, a belief-based model of the auctions is combined with a utility analysis to trade off the certain outcome of the terminating auction against the possible outcomes of the remaining auctions, and hence to place appropriate bids in each.

On the subject of standards, Stefan Poslad from Queen Mary, University of London then presented collaborative work with Phil Buckle and Rob Hadingham of Nortel Networks concerning standards interoperability for multi-agent systems. At present, numerous agencies and agent systems are being developed or portrayed as vehicles for delivering novel types of e-commerce services to users. However, service agents in one agency are probably unable to interoperate or cooperate with agents from another vendor's agencies. Clearly, standardisation in this area would help to create a more ubiquitous market for agent-based services. According to Poslad, standards specifications ought to be grounded within a practical framework that provides a reference implementation, enabling a multitude of developers to build their own implementations. One such open agent platform is called FIPA-OS (FIPA Open Source), originating from Nortel Networks, to promote the uptake of FIPA specifications by agent developers. Poslad described FIPA-OS, which is the first agent platform to be released as open source and is being deployed in several application domains including virtual private network provisioning, distributed meeting scheduling and virtual home environments. FIPA-OS has been demonstrated to interoperate with other heterogeneous FIPA-compliant platforms and is in use in numerous institutes around the world.

Next, John Bigham described a project involving Queen Mary, University of London and Universitat Autonoma de Barcelona that is building a general trust model and security framework for multi-agent systems designed to manage resources in future mobile communications networks. The multi-agent system is being developed to validate the use of agent technology for the control of future mobile communications networks. Bigham described work on possible security problems in the business models of the Shuffle project that provides the context for this work, and outlined trust and negotiation mechanisms to address them. The next step in the work described is to expand the scheme to get a common secure auction model that can be used by general applications.

One of the key problems of recent years in the multi-agent field, long recognised by the UKMAS community, has been the divide between theoretical work in agent-based systems and its practical complement which have, to a large extent, developed along different paths. Presenting collaborative work with Michael Luck, Ronald Ashri (both are at Southampton University) described the Paradigma implementation framework that has been designed with the aim of narrowing this gap. It relies on an extensive formal agent framework, and is implemented using recent advances in Java technology. Specifically, Paradigma uses Jini connectivity technology to enable the creation of on-line communities in support of the development of agent-based systems.

Also in part concerned with the issue of relating theory to practice, Alessio Lomuscio of Imperial College London presented a description of the multi-agent VSK Logic developed with Michael Wooldridge at Liverpool University. The logic enables the representation of what is objectively true of some environment, what of the environment is visible or accessible to individual agents, what these agents actually perceive and, finally, what the agents actually know about the environment. Lomuscio described the logic and its semantics, and showed how it related to agent properties, illustrating his talk with a case study.

Next, Keith Clark described an object-oriented extension of the multi-threaded Qu-Prolog language, developed at Imperial College in collaboration with Robinson. He showed how the language could be used to quickly implement multi-agent applications on the Internet in which agents have both reactive and pro-active behaviours. The different behaviours execute concurrently, as separated threads of an active object that implements the agent. Reasoning agents can thus be defined to perform resource-bounded inference for full first-order predicate logic both to answer questions about what they believe and to check for possible inconsistency before adding new information to a belief store. Such agents can then be elaborated to cooperative reasoning agents that can ask others to engage in sub-proofs on their behalf.

https://doi.org/10.1017/S0269888901000169 Published online by Cambridge University Press

In the final paper session, Jim Doran from Essex University discussed a new, non-standard approach to cognitive processes and multi-agent systems. This approach focuses on the precise, abstract definition of an environmental history that an agent must enable by its actions, on Boolean networks as agent control devices, and on the cognitive processes that these Boolean networks may implement. Luc Moreau from Southampton University presented joint work with Omer Rana from Cardiff University regarding issues in building agent-based computational grids. First the existing infrastructure required to realize the computational grid was discussed, defining such a grid with reference to knowledge and information grids. Then an agent-based approach for the Computational Grid, which is cantered on providing "services" for managing resources, was elaborated. Work reported at the workshop also included the KARO framework, a combined project between Manchester Metropolitan, Manchester and Utrecht Universities. Although there are a number of theories of rational agents that are formulated in combinations of modal logics, the work on practical proof methods for the expressive logics involved in these theories has been sparse. The authors argue that the expressiveness of the KARO framework exceeds that of theories of rational agency and the talk focussed on the current efforts of the project to provide proof methods for the logics used in the KARO framework.

The workshop ended with an interesting and sometimes light-hearted look back at the successes of the multi-agent community in the last five years, which sparked a lively debate between all those present at the workshop.

## 4 Summary

Overall, UKMAS 2000 provided an excellent forum for engaging and lively presentation, debate and discussion within the context of a full programme, but in an informal environment. The success of the workshop in 2000 has led to it being located once again at St Catherine's College, Oxford on the 13th and 14th of December, 2001, co-chaired by Mark d'Inverno of Westminster University and Michael Luck of Southampton University. Further details of the workshop can be found at 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.

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" *The Knowledge Engineering Review* **12**(3) 307–308.

Luck, M, d'Inverno, M, Fisher, M 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" *The Knowledge Engineering Review* **15**(3) 293–301.

https://doi.org/10.1017/S0269888901000169 Published online by Cambridge University Press