

Progress in Multi-Agent Systems Research *

Omer Rana¹, Chris Preist², Michael Luck³

1. Department of Computer Science, University of Wales Cardiff, PO Box 916, Cardiff CF24 3XF, United Kingdom
2. Hewlett Packard Labs, Filton Road, Bristol BS12 6QZ
3. Department of Computer Science, University of Warwick, UK

March 23, 2000

1 Introduction

Continuing the series of workshops begun in 1996 (Luck, 1997; Doran *et al.*, 1997; d’Inverno *et al.*, 1997; Fisher *et al.*, 1997) and continuing in each of the two years since (Luck *et al.*, 1998; Aylett *et al.*, 1998; Binmore *et al.*, 1998; Decker *et al.*, 1999; Beer *et al.*, 1999), the 1999 workshop of the UK Special Interest Group on Multi Agent Systems (UKMAS’99) workshop took place in Bristol in December. Chaired and organised by Chris Preist of Hewlett Packard Laboratories, with support from both HP and BT Laboratories, the workshop brought together a diverse range of participants from the agent community in both the UK and abroad, to discuss and present work spanning all areas of agent research. Although dominated by computer scientists, also present at the meeting were electronics engineers, computational biologists, philosophers, sociologists, statisticians, game-theorists, economists and behavioural scientists, with both academia and industry well represented. Indeed, numbers attending these workshops continue to grow, reflecting the continued and rising interest in agent-based systems. The meeting truly demonstrated the wider view of what the term agency implied to research in other disciplines and the questions raised at the end of presentations were a pertinent reminder of the diversity of the audience.

2 Presentations on the first day

The first keynote speaker, Ed Durfee, presented a wide range of research issues related with planning and co-ordination within agent communities. He discussed co-ordination protocols based on the idea of argumentation between a collection of agents and market-oriented protocols to support co-ordination and control within agent communities. As a way of developing re-usable, extensible, understandable and efficient software, he emphasised the importance of the *negotiation* concept between agents and related the need for this approach within the wider context of software engineering. However, one of the problems with such a negotiation-based approach is deciding what should be modelled within the negotiation paradigm, and this has a great impact on the decisions that would be involved in the process. Durfee also described techniques to reduce the search effort required to reach a decision during a negotiation process, basing it on an auction protocol. He moved on to describe the use of hierarchical planning to find inter-dependencies between collaborating agents, and to divide an

*This report summarises the paper presentations at the Second Workshop of the UK Special Interest Group on Multi-Agent Systems (UKMAS’99).

application being undertaken within an agent community. The hierarchical task network (HTN) was proposed as a way of achieving the latter, and supporting the discovery of agent interdependencies in a decentralised way. This use of HTNs was used as a preamble to introduce the general notion of graph-based planning.

A second aspect presented by Durfee was the use of *congregations* within agent communities, which help dictate the role played by a particular agent. The objective here is to discover congregations (groups, coalitions) that help combine agents offering similar services, requiring similar functionality, or behaving in a similar way. Various approaches to finding such congregations were described, ranging from logic-based approaches to others based on genetic algorithms. A *label*-based approach to the development of congregations was proposed, where agents announce their capabilities to other agents using a predefined label, and supported by specialised ontologies to help others search for agents to form congregations. A pricing structure was proposed as a means to decide the computational costs for joining the congregation, and then operating as a unified entity once inside it. Other negotiation protocols based on neural networks and hierarchic plans were also discussed. To end the talk, Durfee discussed the quality of co-ordination decisions in uncertain environments, and the impact this can have on choosing particular heuristics for planning agent behaviour. He emphasised that to reduce the possible search space for agents, it could be possible to pre-define certain behaviours from which subsequent search could be performed. Various scenarios were demonstrated, both in robot path planning, and as part of the DARPA CoABS project. The objective of the latter was to deduce a possible plan to rescue survivors after an incident, constrained by the fact that only certain routes remained over which a survivor-carrying vehicle could move. Multiple such vehicles had to co-ordinate their plans, especially in a scenario where certain paths were constrained in being utilisable once only. The experiments demonstrated similarities with graph traversal algorithms.

It would be interesting to know the effect of congregations on application performance, and in particular in relation to this ability to group agents based on particular criteria, which may be statically defined (such as a particular service being performed, a particular platform on which the agent service is being run, etc.) or dynamically determined (based on the environment in which the agents operate). For instance, can roles be defined that could help improve application performance? Is it better to have specialised agents that perform specific services rather than more general problem solving agents that can adapt based on application needs? Or, alternatively, is the use of congregations just another way to define agent services at a coarser level of granularity?

In the first of the main paper presentations, Nir Vulkan of the University of Bristol described his work with Chris Preist, and offered an economist's perspective on using agents for predicting stock market trends, discussing the Band-X, Rate Exchange and Min-X criteria for evaluating stock prices. He introduced the *multi-armed bandit problem* used in economics to model stochastic behaviour, where multiple slot machines (each machine being a one-armed bandit) are used as a system for predicting possible pay-offs within a given time interval. Vulkan stressed that when modelling time-series data, as found with stock prices, it is important not to model the time-series directly, but the causes of moves within the data. A market model developed at the Santa Fe institute was used as a problem domain to be predicted by a class of heterogeneous agents operating as a learning classifier system. A number of time series results were presented under different market conditions, and several game theoretic approaches were introduced to model agent interaction, based on reaching a Nash equilibrium within a given number of iterations.

A subsequent talk by Sonia Schulenberg from the University of Edinburgh, describing work with Peter Ross, demonstrated the use of a system for predicting market trends, where each agent is modelled as a genetic algorithm with 100 initial sub-strings representing behaviour (trading) rules. The market itself is modelled as a stochastic process with which agents interact, the market model being

based on real data obtained by the researchers. Within it, each agent is required to operate independently, but is rewarded for a particular action within the market. The genetic algorithm is used to maintain those sub-strings that generate a greater improvement in the pay-off received by the agent. Genetic algorithm operators such as mutation and crossovers are restricted to operate on sub-strings within each agent. The work presented an interesting perspective on using evolutionary computing, rather than logic based approaches to modelling agent behaviour. A number of results of the system were presented, including experiments where the market model was changed in unpredictable ways.

The next two talks described progress on threads of research that had continued since being presented at previous UKMAS workshops. First, Rogier van Eijk of Utrecht University presented his work with de Boer, van der Hoek and Meyer on developing operational semantics for agent communication languages (ACLs), where the need for developing well understood semantics for languages such as FIPA ACL and KQML was emphasised. An approach based on CSP was introduced, with message interchange semantics expressed using CSP syntax. Van Eijk stressed that their work was not restricted to any one agent communication language, and indicated that this exercise was a necessary first step towards creating a programming language for agents.

Then, Simon Parsons of Queen Mary and Westfield College described his work with Sabater, Sierra and Jennings on developing agent systems made up of *multi-context* systems, where each agent provides a specialised role within a federated environment. The system is composed of *units* that provide theories on particular agent behaviours, and *bridging rules*, which combine these units together. An approach to negotiation based on argumentation using *rebutting and undercutting* was presented, which could aid agents to specialise their particular roles or behaviours. One objective is to compose agent systems in a hierarchical manner, using bridging rules, that could facilitate the creation of agents for use in multiple applications — such as a planner, a resource manager, a goal manager, a social manager, a rebutting agent, and an undercutting agent. The part of the system currently implemented was demonstrated, in which a GUI could be used to put together different rules based on inheritance properties in object-oriented systems, but further abstractions are necessary to develop more functionality. A Java-based system was described in which multiple Java Virtual Machines were connected together, each of which could employ a multi-threaded Java architecture internally.

Prior to a discussion on agents in spacetime involving Jim Doran and Mike Wooldridge, Marco Arranz from Manchester Metropolitan University provided a constraint-based approach to combine plans from multiple agents. The aim of this work involves establishing a framework within which plans of multiple interacting agents could be reconciled, to enable cooperative relationships and sharing between agents within a community. A *coordination* algorithm was proposed, which involves agents exchanging a *transferable* part of their plan followed by a negotiation phase when a potential mutually useful plan is detected. The long term objective of this work is to study the impact of plan-based relationships between agents, and how this could affect the internal behaviour of each member of the agent community. This work is an automata-based approach to negotiation, and is therefore very much a remnant of path-planning approaches from robotics. The paper highlighted an important area of decentralised management between collaborating agents, and the framework being developed could find wide use in multi-agent communities.

3 Discussion on Security and Performance

Towards the end of the first day, there was a discussion session on security and performance issues in multi-agent systems. Stefan Kirn introduced the discussion on security issues, and identified a layered approach to handling the problem. The layered approach is based on different ways in which agent

systems could be affected within an organisation, and includes aspects of wire tapping, altering of data, copying and replaying, denial of service, and abrogation and masking. He emphasised the importance of creating a framework within which security decisions could be evaluated, and how this approach could be used to maintain and evolve security concerns within an enterprise. Subsequently, Stefan also introduced the new German initiative in agent systems for business problems, and discussed two areas of interest: health care and enterprise management.

Omer Rana of the University of Wales raised the importance of performance engineering within agent communities, which he argued must not be ignored, under the false premise that we should wait for agent research to mature before these issues become relevant. Performance can be viewed from many perspectives, and is not just about how fast an application runs. Issues that extend beyond *platform* performance, and consider the wider notion of *application* performance, should be evaluated in the context of agent communities. We need to identify issues which are specifically agent oriented, to guide the notion of performance engineering within agent systems, and these therefore must be studied together with software engineering approaches — perhaps, focusing on themes such as Conversation Policies, creating Congregations/Groups within a multi-agent domain, Roles and Policies within an agent community etc — and consider the effect on performance in making some of these decisions. Performance should therefore be related to the wider issues of reliability, efficiency, fault tolerance, specialisation, heterogeneity, in addition to the traditional aspects of execution time. This is particularly important when considering special needs imposed by emerging applications in mobile and distributed computing, involving a diverse range of different devices, offering services that can vary in time and space, and which involve collaborating agents co-operating and competing within dynamic networks. The notion of match making and service discovery within such environments must be connected with scalability and performance — and approaches to develop agent communities must handle these directly. Existing themes explored by d’Inverno and Luck, for instance, based on Z notation (for instance), should be augmented with performance models. In a similar way, as a community, we need to define benchmarks and metrics that could help us compare and contrast different agent systems, to support software engineering themes within agent systems.

4 Presentations on the second day

The second day began with an invited presentation from Owen Holland of Cyberlife. Holland presented a wide range of issues in developing agent communities that interact with each other based on evolutionary paradigms. A number of approaches were introduced based on the concept of *emergent behaviour*, whereby the combined behaviour of a collection of entities leads to some useful behaviour that was not initially predicted. The developer specifies local behaviours for each entity, and does not know, or cannot define, the combined behaviour of the community of such entities. This approach to developing agent communities is derived from similarities with social insects, and generally referred to as *swarm intelligence*. Holland presented a number of examples of such systems based on ant colonies, and the collective behaviour in a beehive. Cue-based action, whereby a group of ants undertake specialised pre-defined operations and perform workload sharing, based on cues from other ants, is a paradigm that could be implemented in developing artificial computing systems. Other working patterns in ant colonies are also possible, such as locating and managing a food source, and creating clusters of objects within the nest (later demonstrated with a robot example using pucks).

In this context, the use of *sematectonic stigmergy* can be adapted from computational biology for use in managing computer networks. For example, a network management system based on this approach has been developed with Hewlett Packard Labs, in which a collection of ants model work

packets that can be migrated within a network based on utilisation of network links. The intention here is to avoid the costly computation of routing tables between nodes within a network, and enable alternative paradigms based on computational biology to automatically deduce shortest paths between interacting nodes. According to Holland, the system was a huge success, and demonstrated the use of collective behaviour within a colony of interacting agents (or migrating ants) for solving a very hard optimisation problem. The speaker touched on a number of issues that could be borrowed from computational biology and, in particular, socially interacting insects, such as:

- collective goal sequencing;
- providing sequencing and co-ordination between groups of tasks undertaken by such agents;
- recruitment and competition, especially with reference to ant colonies, where army ants can be recruited to help worker ants if special patterns of work (behaviour) are encountered;
- extending the previous notion to foraging for work, where workers without work can ask for additional tasks to be performed from those that are otherwise occupied (a technique for achieving automatic load balancing across dynamic roles); and
- work adjustment undertaken by participants within the community, based on migrating work packets upstream or downstream depending on the availability of particular workers within the pool.

Holland presented a fascinating video of interacting robots that were required to pile pucks based on their colour. It was found that over a period of time, a clustering behaviour emerged automatically, which was particularly noteworthy, as there was no explicit algorithm in operation causing this effect. Adjusting the conditions within the environment, such as slope of the field within which the robots were operating, floor friction etc, led to variations in the number of puck clusters created — and Holland subsequently offered a “kama sutra” of pucks within the robot world. The general lesson from these experiments, according to the speaker, was the fact that simple actions performed by these robots can lead to unexpected behaviours that could not have been predicted or, in some instances, directly programmed.

An alternative approach, based on a global behaviour rather than localised interactions in interacting robots was also presented, resulting from the use of cellular automata to plan paths through a maze. The notion of using specialised *fluid* motions, whereby a chemical ‘goo’ could be placed within a maze, and which could concurrently explore various routes through it, similar to the outward expansion of a spilled fluid or gas, was an approach to exploring alternative plans. In this instance, from an initial starting point in the middle of the grid, the chemical would spread out and, in the process, automatically explore possible paths available within the system. Result from a number of such experiments were described, with these latter experiments being based on a *continuous* computation, as opposed to the earlier experiments with interacting robots, which were necessarily discrete. Finally, Holland also briefly touched upon the emerging research in consciousness, and how it applies to agent communities.

Holland’s talk provided an interesting insight into the alternative approach to developing agent communities, undertaken traditionally within the A-Life research community. The notion of useful behaviour from a collection of simple ones is an appealing concept, particularly as it seems to uncover a problem-solving approach employed by nature. However, the time to reach a useful behaviour within such a community of self-interested, and often simple participants, is an open question. For example, in the demonstrations shown, the clustering behaviour within the robot community took a long time

to achieve — and a question that could be posed to supporters of such methodologies is whether this approach can be generalised to achieve problem solving within other software environments, within guaranteed time frames. On the other hand, the discovery of emergent properties within a community of interacting agents is often frowned upon by the logic-centric agents researchers, who suggest that it is to be avoided rather than encouraged. From this viewpoint, unspecified emergent behaviour is necessarily bad, and therefore should never happen. Can this view be altered to consider a middle ground, whereby emergent behaviour could be constrained in a useful way within a logical framework, so that the best of both approaches are utilised? Do emergent properties give us ways of handling scenarios that cannot be easily (or efficiently) described within a logic framework?

The first of the paper presentations on the second day began with Ana Bazzan of UFRGS describing work with Wahle and Klugl on the use of agents in traffic modelling. By modelling drivers as social agents in a multi-agent system, allowing their behaviour to be investigated and predicted, they have implemented a prototype traffic simulation system. There are two levels of analysis: at the tactical level, the sub-cognitive aspects of traffic and driver action are considered, while at the strategic level, more deliberative models are used that include mental states like emotions, preferences and intentions.

In his paper, Beer was concerned with the delivery of effective integrated community care with the aid of agents. Together with colleagues Bench-Capon and Sixsmith at the University of Liverpool, he tackled the problem domain of patient care in the community by developing an agent architecture for home monitoring community alarms, care management systems, and so on, in which each of these components is represented by an agency. Using wrappers and conversation classes, for example, to map to existing systems, an effective prototype has been developed.

The last of the morning papers was presented by Cao of the University of Warwick, who described an agent-based modelling and simulation approach to overcoming the difficulties associated with evaluating advertising effectiveness. A demonstration version of a system to provide a Simulation Environment of Advertising (SEA) has been developed to study the effectiveness of banner advertising. It should allow users to construct their own models from different components, including consumer model selection, negotiation modeling and analysis modeling. Though the work is still early, it points to ways in which agent-based systems may be used for advertising research.

After lunch on the second day, a series of short presentations followed. The first, by Tony Kakoudakis from Manchester Metropolitan University describing work with Michael Fisher, involved the specification of what constituted a *group* in the context of Concurrent METATEM. This language has been widely used in the context of executable logics in multi-agent systems, and enables both the specification of the behaviour of an agent and its subsequent execution. The approach in the talk made use of temporal logic to specify agent behaviour, subsequently translated to Separated Normal Form, which was then used to group agents, restricting interactions between agents to be within the group. The approach was generalised so that the notion of *groups* can be viewed at multiple levels of granularity, leading to the observation that *agents* and *groups* could refer to the same entities — indicating interaction to be intra-group (between agents) and inter-group (between groups). The recursive nature of the approach — where agents and groups refer to the same entity in the limiting case — is a very interesting aspect of this work. However, the emphasis in the framework is more on how to add and remove members from the group, and its subsequent management, rather than (and presumably after) the identification of group members using suitable metrics.

The second short presentation came from Eduardo Alonso from the University of York, whose work with Kudenko used a conflict-resolution based approach to demonstrate multi-agent reinforcement learning. The particular scenario presented included a military application involving a battle-tank moving on a structured grid — clearly suggesting the commonality of this work with that of

robot path planning, albeit under uncertainty. Their work makes use of inductive logic programming and explanation-based learning to build learning agents, while employing background knowledge for assisting planning and learning in a multi-agent system, and logic for knowledge representation.

The next presentation was delivered by Professor Peter Gray of Aberdeen University on the KRAFT project, involving Aberdeen, Liverpool and Cardiff Universities and BT. The project involves a collaboration of agents that wrap legacy databases and other information sources, for facilitating knowledge sharing and re-use. With specialised agents such as mediators and facilitators undertaking roles to assist knowledge users and managers, the emphasis in the project is on managing and handling constraints in the context of knowledge sharing. The presentation demonstrated the use of agent technology in an area that could impact emerging areas such as electronic commerce and information management, and a comparison with a related US project, InfoSleuth (at MCC)¹ was also given.

Rafael Bordini of UFRGS Brazil ended the final paper presentation session with a thought-provoking talk about his work with Campbell and Vieira, using ontologies to share knowledge between collaborating agents. A formalism was presented, which supported taxonomic relationships to be derived from ontologies defined using this formalism, and used Luck and d’Inverno’s formal agent framework (in Z) as a basis for their description. Relating his work to ethnographic studies in social anthropology, Bordini exemplified the diversity of agent research, and claimed that ontologies are not designed by programmers or designers of multi-agent systems, but that agents may change them evolutionally. This work is therefore aimed at enabling agents to develop dynamic ontologies based on descriptions that can be inferred within a society of interacting agents — a major undertaking by any standard.

5 Summary

A number of approaches to agent communities were discussed on the first day of the workshop, ranging from planning methods for agents within a community, to market based methods for reaching convergence (equilibrium) within a collection of problem solving agents. It was interesting to note that planning decisions currently being considered are driven by situated robotics, with influence from research in path-planning and obstacle avoidance. The extension of these approaches to software agents which need to operate within an environment composed of databases, workstation clusters, specialised software with licensing constraints etc, is not clear.

The use of compositional approaches to agent based systems was also interesting, especially if this technique can be used to hierarchically construct agent systems — similar to the object-oriented approach, whereby some base (default) functionality could be extended (overridden) to describe specialised behaviours. A core set of features could be abstracted for particular domains, which could then be extended by a programmer. Core features could be a *negotiator* for instance, and specialised roles could be an *undercutter* etc. Agent interaction is often ignored, and the use of bridging rules between agents is also an interesting concept, not dissimilar to the *connector* idea introduced in the ADE work at Potsdam (Horn et al, 1999). Connecting agents which can perform specialised roles, and identifying communities or congregations of such agents could be useful in constructing large scale agent communities. Although not obvious, the static allocation of agents to communities is one way to create such congregations. However, automatic allocation of agents to such communities seems an interesting way to create specialised co-operative and competing communities that could provide generalised problem solving within specialised domains.

Such approaches to software agent design are certainly useful, as they facilitate the current emphasis on distributed, de-coupled systems. A design approach that builds on current efforts within the

¹<http://www.mcc.com/projects/infosleuth/>

object-oriented community, for instance, would be useful, as they would leverage the existing investment in object and component technologies. Perhaps there is a need to unify some of these tools into a methodology for software engineering agent systems that could benefit from the techniques outlined in the workshop — such as the use of bridging rules and roles. In the same vein, developing semantics for agent communication languages is also an important research goal, and the paper presented on the first day provided a useful perspective on this problem.

The workshop brought together a range of different issues and applications in single and multi-agent systems. Two aspects lacking in the workshop were user-agents (such as graphical interface agents based on multimedia characters in virtual worlds) and mobility aspects associated with agents. The workshop did however highlight the importance of looking at a range of techniques that extend beyond logic-programming, for modelling agent systems. The evolutionary approaches often employed could prove useful in scenarios where a mathematical or logic oriented approach was not possible — and the integration of these two approaches must be investigated further.

Indeed, the importance of a multi-disciplinary perspective to agent communities, by which it is essential to import ideas from a range of different disciplines into agent research and not restrict it purely to computer science or economic-based approaches, which seem to be dominant at present, was clearly highlighted at the workshop. The importance of taking account of particular problem domains was also emphasised, and here the close interaction of industry and academia was necessary to identify useful business models and problem domains which can benefit from agent technology. This is one of the key aims of the UKMAS workshops, to foster and promote collaboration among industrialists and academics, and to contribute to the development of the field. The evidence from UKMAS'99 on that score is positive.

Following on from the success of this event, and the previous events, the next UKMAS workshop will be held in Oxford on the 14th and 15th of December 2000, chaired and organised by Mark d'Inverno of the University of Westminster. Details of the workshop are available at <http://www.wmin.ac.uk/~dinverm/ukmas2000.html>.

References

- Aylett, R, Brazier, F, Jennings, N, Luck, M, Preist, C and Nwana, H, 1998. "Agent Systems and Applications" *The Knowledge Engineering Review*, 13(3): 303–308.
- Beer, M, d'Inverno, M, Jennings, N, Luck, M, Preist, C and Schroeder, M, 1999. "Argumentation and Negotiation" *The Knowledge Engineering Review*, 14(3): 285–289.
- Binmore, K, Castelfranchi, C, Doran, J and Wooldridge, M, 1998. "Rationality in Multi-Agent Systems" *The Knowledge Engineering Review*, 13(3): 309–314.
- Decker, K, Fisher, M, Luck, M, Tennenholtz and UKMAS'98 Contributors, 1999. "Continuing Research in Multi-Agent Systems" *The 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" *The Knowledge Engineering Review*, 12(3): 315–321.
- Doran, JE, Franklin, S, Jennings NR, and Norman, TJ, 1997. "On Cooperation in Multi-Agent Systems" *The Knowledge Engineering Review*, 12(3): 309–314.
- Fisher, M, Müller, J, Schroeder, M, Staniford, G and Wagner, G, 1997. "Methodological Foundations for Agent-Based Systems" *The Knowledge Engineering Review*, 12(3): 323–329.

Horn, E, Kupries, M and Reinke, T, 1999. “Properties and Models of Software Agents and Prefabrication for Agent Application Systems”, Hawaii International Conference on System Sciences (HICSS-32), Software Technology Track.

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” *The Knowledge Engineering Review*, 13(3): 297–302.