

Foundations of multi-agent systems: issues and directions

MICHAEL LUCK

Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK

The last ten years have seen a marked increase of interest in agent-oriented technology, in several areas of computer science, including both software engineering and artificial intelligence. Agents are being used, and touted, for applications as diverse as personalised information management, electronic commerce, interface design, computer games, and management of complex commercial and industrial processes. Several deployed systems already use agent technology, and many more commercial and industrial applications are in development.

Though there is, as yet, no agreement as to exactly what constitutes an agent, an agent can be viewed as a self-contained program capable of controlling its own decision-making and acting, based on its perception of its environment, in pursuit of one or more objectives. Multi-agent systems are computational systems in which several agents cooperate to achieve some task which might not, or not as easily, otherwise be achieved.

It is a relatively new field of investigation, but one which is distinct yet widely applicable, and which promises to have a significant impact in a large number of areas. However, the relative immaturity of the field does provide some resistance to the speedy uptake of agent technology. The lack of consensus over terms and concepts, for example, is one such constraint, but other concerns relate to the need for a detailed understanding of the mechanisms and structures that must be involved in, and underlie, both single-agent and multi-agent applications, and for the supporting tools and techniques that are required for their development.

One way to address these problems is to focus efforts on key aspects of agents and multi-agent systems and so develop a solid foundation on which to base practical systems. There are many such issues spanning a diverse range including, for example, those relating to the nature and roles of autonomy, communication, cooperation, coordination, negotiation, reactivity, pro-activeness, etc., development tools and techniques, methodologies, formalisms, and so on. This list is not exhaustive, merely indicative of the kinds of issues that are of interest.

The *First UK Workshop on Foundations of Multi-Agent Systems*, held at the University of Warwick in October 1996, was intended to encourage and support activity in the research and development of multi-agent systems, in both academia and industry and to provide a forum for the discussion of exactly the issues listed above. The rationale for this aim is that multi-agent systems are widely expected to become a significant technology in the 21st Century, and it is crucial that both academics and industrialists within the UK have access to a forum at which current research and application issues are presented and discussed.

The workshop comprised three panel discussions on cooperation, formalisms for multi-agent systems, and methodological foundations. Each of the panel members was asked to respond to a set of questions on the panel subject in advance of the workshop, to be used as the starting point for the discussion.

Cooperation, the fundamental relationship between individual agents in any multi-agent system, is rich and varied in both terminology and concept, and for that reason is still difficult to circumscribe. The first panel was aimed at exploring the nature of cooperation and its connection

with other forms of interaction between agents, the different types of cooperation, and the ways in which it is initiated and how it is sustained.

The second panel was concerned with the use of formal methods and techniques in the field of multi-agent systems. There is a growing body of work reported in the literature that concentrates on theoretical aspects of agents and multi-agent systems, and a complementary body of work concerned with building practical systems. However, the two have typically been unrelated. This panel focussed on the role of formalisms in multi-agent systems and their relation to the development of practical systems, on the specific qualities required of such formalisms by the multi-agent domain, and on the relative merits of several current formalisms and their application to multi-agent systems.

Finally, the methodological foundations panel addressed perhaps the most immediately critical issue of how to set about the development of multi-agent systems. If we are to be able to exploit the benefits offered by these systems on a large scale, we need principled methods for development. Though this has much in common with the development of other complex software, the particular components required in agent-based systems, coupled with the rapid progress of the field, suggest that this is an important area for investigation. The panel considered both general development methods, and current approaches to the problem.

Around forty people attended the workshop, including participants from Germany and the United States, and with strong representation from industry. This cross-section of the agent-based systems community, and the panel format of the workshop fostered a focussed and effective interchange of ideas that made for a lively and informative day, and that is reflected in the summaries of the panels that follow. After the success of this first workshop, a second workshop is scheduled for December 1997, again to be held at Warwick. Details both of the organising Special Interest Group and of the workshops can be found on the FoMAS web pages (<http://www.dcs.warwick.ac.uk/~fomas/>).