The Knowledge Engineering Review, Vol. 14:2, 1999, 291–294
Printed in the United Kingdom

Making money from agents

DAVE CLIFF[1], JÖRG MÜLLER[2], DIVINE NDUMU[3], CHRIS PREIST[1] and MIKE WOOLDRIDGE[4]

[1]*Hewlett-Packard Laboratories, Filton Road, Stoke Gifford, Bristol BS34 8QZ, UK. Email: cliffd@hplb.hpl.hp.com*
[2]*J. P. Wiley and Sons Ltd, International House, Ealing, London W5 5DB, UK. Email: jpm@wis-dev.wiley.co.uk*
[3]*BT Laboratories, Martlesham Heath, Ipswich IP5 3RE, UK. Email: ndumudt@info.bt.co.uk*
[4]*Department of Electronic Engineering, Queen Mary and Westfield College, University of London, E1 4NS, UK. Email: m.j.wooldridge@qmw.ac.uk*

The discipline of agent technology is becoming a mature research area, both from a scientific and an engineering perspective. Strong theoretical and empirical analysis, together with practical experience of deploying systems to solve real problems, give this community the opportunity to make an impact in the world of business. At UKMAS'98, a panel of academics and industrialists met to discuss exactly how this could be done; how could agents and multi-agent systems be used to actually make money? The panellists considered the following questions;

1. Some technologies, such as the spreadsheet, allow you to do the same thing as before, but more easily. Other technologies, such as virtual reality, allow you to do things that were impossible before. What previously impossible things will agent technology enable? Or what difficult things will become easy?
2. Who will find these changes useful or desirable? Will they be prepared to pay for such changes?
3. What kind of company will be in a good position to sell these changes?
4. What research will be necessary to bring agent technology to this point?

What previously impossible things will agent technology enable? Or what difficult things will become easy?

As with any discussion of agent technology, different members of the panel had different views of what an agent or multi agent system consisted of. Wooldridge identifies *agents in the small* and *agents in the large*. He defines agents in the small as systems that make use of the agent concept, but which do not necessarily make use of techniques that agent researchers would regard as agent technology, and cites the Microsoft office assistant as an example. He defines agents in the large as coarse-grained computation entities, each permanently containing at least one thread of control, typically endowed with some sort of reasoning, planning, or learning capability. Ndumu similarly divides agent technology into two levels, roughly corresponding to agents in the small and agents in the large. He defines the agent approach as a way of viewing systems at these different levels; at the individual agent level, this view is of *delegated responsibility*, while at the multi-agent systems level it is one of *situated knowledge level interaction between autonomous systems*. Cliff takes an artificial life perspective, considering agent communities consisting of *animats* and possibly people. He defines animats as artefacts that have animal-like capabilities of coordinating sensory input with motor output for survival and adaptation in their particular environments where these environments are not known in advance, are unpredictable, and are unforgiving of mistakes. Müller takes a general view of agent technology, defining an agent as a program that is smart, and can do things in a way that a human would consider intelligent.

Wooldridge points out that agents in the small are already being deployed by companies like

Microsoft to make desktop applications easier to use. He observes that the competitive advantage for desktop applications like Microsoft's Office does not lie in the raw functionality of the application, but in the ease with which this functionality can be exploited by a user, and that agents in the small can help with this. He does warn, however, that it is not currently clear that greater improvements could not be made by more "traditional" refinements to system user interfaces. Ndumu believes that the main benefit of agents in the small will be the provision of personalised value added behaviour, allowing delegation of tasks to agents. In addition to personalisation, Müller adds two additional functions of agent technology: multi-source pattern recognition and matching, and dynamic distributed resource allocation. Based on these functions, he identifies four likely application areas for agent technology:

- Enriching personal communication.
- Personalised information retrieval from public and corporate information sources.
- Intelligent data mining applications for sales and marketing.
- Intelligent support of automated e-commerce transactions, auctions and resource discovery applications.

Wooldridge and Ndumu agree that multi-agent systems provide a new perspective on conventional software engineering, so will not be able to do anything fundamentally new. However, this perspective allows certain problems to be tackled more easily. Ndumu identifies enhanced systems interoperability as the key. The meaningful interaction between heterogeneous and distributed systems this enables could, he feels, have a large impact on computing, business and society as a whole. Wooldridge observes that multi-agent systems explicitly address the problem of how to make decisions in a society composed of entities that do not share the same goals or intentions. Many important enviromnents—including, perhaps most importantly, the Internet—belong to this category, so can benefit from applications developed in this way.

Cliff observes that animats with the full cognitive and behavioural sophistication of an insect, let alone a human, will require a large amount of further research. However, even relatively simple animats can be programmed to exhibit *apparently* sophisticated behaviours, and to interact with each other (and/or humans) in realistic ways. Hence, realistic communities of agents can be created, and can interact with people.

Who will find these changes useful or desirable? Will they be prepared to pay for such changes?

Müller identifies two main beneficiaries of agent technology; the *knowledge worker* and the *virtual organisation*. A knowledge worker is someone who works while embedded in the information processes of an institution, enterprise or community. The knowledge worker requires tools to support location-independent and intelligent services for information access, information transfer and process management. A virtual organisation is a complex, open, dynamic structure of individuals and organisations which come together to provide a product or service. It requires IT systems that are able to support change of processes and organisational structures, business transactions and information retrieval and discovery across the organisation. The functions of agents that Müller identified in the previous section could benefit both the knowledge worker and the virtual organisation. However, Müller cautions that such technology will only be taken up if the need for action is great, the potential savings are large enough, and the company/individual has a large enough IT budget to make the investment. Additionally, Wooldridge warns that the takeup of agents in the small will only occur if the agent-based approach really does provide advantages over existing techniques for solving such problems.

Turning to agents in the large, Wooldridge observes that anyone who wants to engineer systems that apparently contain autonomous software components stands to profit from multi-agent systems research. Ndumu argues that the increasing service-based nature of the economy will mean multi-agent technology is increasingly important; the service sector requires the provision/

application of relevant, timely and situation specific knowledge to time-critical problems. This requires the integration and manipulation of resources from many and disparate sources, something multi-agent systems are well able to do.

However, Ndumu sounds a note of caution. Adopting multi-agent technology will require a major paradigm shift in the IT industry and a change of business practice in industry in general. For the agent approach to make any significant impact on a company, then it must reengineer most of its business processes, and convince other players in the industry to do the same. Because of this, few companies are prepared to invest in this approach. Ndumu predicts that early adopters of agent technology will be small startups, and companies under threat of loss of market share.

Cliff observes that, as with previous innovations in computer science, that some of the first people to make money out of agents will be the snake-oil salesmen. They will talk up the technology and build impossible expectations in the naive (while charging hefty per-day consultancy fees) and walk away before anyone realises the new medicine doesn't work any better than the last batch.

On a more positive note, Cliff argues that the apparently sophisticated behaviour of animats makes them ideal for use in industries where people pay good money to be fooled; namely entertainment industries, such as computer games and films. Already, flocking and herding scenes in some movies are generated using simple animat technology, and Furbies are simple autonomous agents. Cliff observes that computer games involve interaction with agents, though currently most of these are very dumb indeed. Games designers often get around this by overpopulating the game with these dumb agents, making them aggressive and fast, so the issue then is whether you have sufficient hand-eye coordination to kill all the agents before they kill you. However, while this appeals to the adolescent male, many people would prefer to outwit a few intelligent opponents, not fill idiot hordes full of lead. Animats, providing seemingly intelligent group dynamics, can therefore lead to more stimulating games. Similarly, in computer animated movies, autonomous software agents can be used to act as "extras" and "bit part players". Cliff also predicts the merging of the computer game and movie industries, creating truly immersive "interactive movies".

What kind of company will be in a good position to sell these changes?

Müller observes that companies that are very close to the basic software/communications infrastructure and content usually sell value-added technology most effectively. For that reason, he predicts that the first companies to make money from agents will be in the media, telecoms, publishing, IT and entertainment sectors. In addition, small startups are likely to fill niches left by the big players.

Wooldridge sounds a note of caution with respect to the commercialisation of multi-agent systems. If one considers how long it has taken objects to become adopted by the mainstream, then it could clearly be decades before agents are also accepted as such. Just as it is difficult for us to accurately quantify the benefits that object-oriented solutions bring, so it will be difficult to precisely determine the benefits of an agent-oriented solution. However, the benefits of agents in the small are more self-evident, so any software company that wants to make its applications easier to use might want to exploit agent technology.

Ndumu is also cautious about the uptake of multi-agent systems by businesses, this time because of the structure of the industry. A middleware infrastructure of intraoperability enhancements, together with agent-based applications exploiting this middleware, will be necessary. However, different companies are responsible for different parts of this application stack, creating a potential vicious circle; agent-based systems can't be developed without the appropriate middleware, but middleware companies are reluctant to invest without an existing established market. Partnerships may be the way forward. Alternatively, a major player could use its market power to force its contractors/suppliers to adopt an agent-based solution.

In the previous section, Cliff identified that the entertainments industry will be significant users of agent technology, both in games and films. However, he does not feel that they will necessarily develop this technology. So long as appropriate APIs are agreed, companies skilled in creating

believable agents for entertainments could sell instances of their agent control engines repeatedly to multiple buyers, tailoring each unit sold to the demands of that particular role. Such companies could be small, quick moving startups.

What research will be necessary to bring agent technology to this point?

With respect to agents in the small, Wooldridge feels that the main research issues are understanding how to make unobtrusive but helpful agents, that can actively assist users without being stupid, stubborn, or overbearing. With respect to agents in the large, he identifies two crucial issues currently addressed by the multi-agent systems community:

- Understanding agents as a software engineering paradigm—how do they fit into the software engineer's world-view? Exactly what characteristics are appropriate for a domain to be viewed as suitable for an agent solution?
- What algorithms and techniques can be used to enable agents with possibly conflicting goals, intentions, etc., to cooperate and negotiate effectively?

Ndumu identifies the main technical drawback to agent deployment as lack of standards for agent interoperability. From an industrial viewpoint, he feels that more effort should be directed at building and deploying practical agent applications. This is important for three reasons:

1. To imbue industry with confidence in the technology.
2. To create a substantial base of computer systems for new ones to interoperate with.
3. To highlight the particular agent issues that would benefit from further research.

From a research perspective, he feels that agent responsibility and ease of human–agent interaction are key issues.

Conclusions

The discussion focussed roughly around three types of agent:

- The personal assistant agent (or agent in the small) which participants felt was already being used and generating income for some companies, and could play an increasingly large role in the life of the knowledge worker and virtual enterprise.
- More sophisticated multi-agent systems are further away from fulfilling their potential role in industry. This is partly because new technologies take time to be accepted and partly because technical and business challenges, such as the need for interoperability, must still be overcome.
- Communities of agents behaving in believable ways are already used in the entertainments industry in simple ways. There is great scope for more sophisticated agents of this form to be used in movies and games, possibly even generating a new genre of interactive movie.