Software Agents: An Overview

6 Some General Issues and the Future of Agents

"Smart agents are here to stay. Once unleashed, technologies do not disappear" (Norman, 1994, p. 71).

We agree fully with these sentiments which is why we believe agent technology is not a passing fad. We have now overviewed a broad range of work which goes under the banner of 'agents'. We have outlined their various promises as well as their challenges. However, apart from such technical issues, there is also a range of societal (i.e. social) and ethical problems which are looming. Donald Norman writes:

"Probably all the major software manufacturers are exploring the use of intelligent agents. Myths, promises, and reality are all colliding. But the main difficulties I foresee are social, not technical. How will intelligent agents interact with people and perhaps more importantly, how might people think about agents?" (Norman, 1994, p. 68).

We disagree with Norman as regards the major technical hurdles ahead; as shown in the previous section, there are some extremely demanding technical issues to be resolved in most of the agent classes reviewed. However, we agree that he poses an extremely pertinent question. There are issues which society would have to grapple with through various legislations and they would be very thorny. They include the following:

- Privacy: how do you ensure your agents maintain your much needed privacy when acting on your behalf?
- Responsibility which goes with relinquished authority: when you relinquish some of your responsibility to software agent(s) (as you would do implicitly), be aware of the authority that is being transferred to it/them. How would you like to come back home after a long hard day at work being the proud owner of a used car negotiated and bought for, courtesy of one of your (Kasbah) software agents? How do you ensure the agent does not run amok and run up a huge credit card bill on your behalf?
- Legal issues: following on from the latter, imagine your agent (which you probably bought off-the-shelf and customised) offers some bad advice to other peer agents resulting in liabilities to other people, who is responsible? The company who wrote the agent? You who customised it? Both? We envisage a new raft of legislation would need to be developed in the future to cover software agents.
- Ethical issues: these would also need to be considered. Already, David Eichmann (1994b) is already concerned enough about the ethics of software agents that he has proposed an agent etiquette for information service and user agents as they gather information on the WWW. They include the following:
    - Agents must identify themselves;
    - They must moderate the pace and frequency of their requests to some server;
    - They must limit their searches to appropriate servers;
    - They must share information with others;
    - They must respect the authority placed on them by server operators;
    - An agent's services much be accurate and up-to-date.
- Etzioni & Weld (1994) have proposed others including:
    - Safety - the agent should not destructively alter the world;
    - Tidiness - the agent should leave the world as it found it;

- Thrift - the agent should limit its consumption of scarce resources;
- Vigilance - the agent should not allow client actions with unanticipated results.

However, such issues are not that critical immediately, but would be in the medium to long terms. In the short term, we expect some basic agent-based software to be rolled out e.g. some basic interface agents such as mail filtering or calendar scheduling agents. More basic mobile agent services would also be provided in the short term. We can also predict comfortably that many vendors would claim that their products are agent-based even though they most certainly are not. For example, we are already hearing of 'compression agents' and 'system agents' when 'disk compressers' and 'operating systems' would do respectively, and have done in the past. As Guilfoyle (1995) warns:

"there is a danger, however, that customers may be disappointed by the gap between colourful press reports about smart agents handling half the work, and the reality of 'if-then' rules for message routing", p. 139.

In the medium term (3 to 5 years), some more decent agent applications would be rolled out for most of the classes of agents overviewed in this paper. Perhaps, there would be collaborative agents and/or integrated heterogeneous agent applications doing limited, but *real* workflow or air traffic management or controlling real telecommunication networks say, etc., rather than just simulations. Useful, but limited, interface agents should be available which perform roles including the following: eager assistants, WWW guides, memory aids, entertainment and WWW filters/critics. Indeed, the HOMR/Ringo system is already being marketed as Firefly by Agents Inc. - a Boston-based agents company. Open Sesame (Caglayan *et al.*, 1996) - another interface agent which employs a neural network learning technique, is already on the market from Charles River Analytics (Cambridge, MA). More mobile and information agent applications and languages would soon be rolled out by vendors. We expect reactive and/or hybrid agent technology to start delivering some real everyday industrial applications during this period. Furthermore, during this medium term, we expect the WWW to be commercialised, to some degree, enabling agents of different classes to play a role in paying for services and performing some restricted buying and selling on our behalf, as Kasbah agents propose to do. We would also start seeing a proliferation of specialist agents conferences: agents in the aviation industry, agents in law, etc. We also envisage that the new domain of agent-based software engineering will grow from strength to strength. In the long term (> 7 years), we would expect to see agents which *approximate* true 'smartness' in that they can collaborate and learn, in addition to being autonomous in their settings. They would possess rich negotiation skills and some may demonstrate what may be referred to, arguably, as 'emotions'. We caution against the usage of such words as the latter, not least because the agent literature abounds with such vocabulary which have subtle and complex meanings in the human context (Smith, 1996a). However, it is also at this stage society would need to begin to confront some of the legal and ethical issues which are bound to follow the large scale fielding of agent technology. In the long term too, agents would also provide another design approach to constructing complex pieces of software.

Indermaur (1994) writes:

"they can't fly yet, but intelligent agents and smart software are beginning to walk" p. 97.

Our view is different: using Indermaur's metaphor, we argue that limited software agents are just to about to 'crawl' out of research laboratories; apart from a few interface and information agents, there is still a very long way, yet, to them 'walking', yet alone to their adolescence.

In summary, we believe, like Greif (1994), that agents can have an enormous effect, but that this will appear in everyday products as an *evolutionary* process, rather than a *revolutionary* one predicted by much press coverage. For example, the Ovum figures quoted in Guilfoyle's (1995) paper suggest a seven to eight fold increase in the agents market, in Europe and USA alone by the year 2000 (i.e. in four years time) - this suggests a revolution! Greif notes correctly that agents would initially leverage simpler technologies available in most applications (e.g. word processors, spreadsheets or knowledge-based systems). Then agents would gradually be evolved into more complicated applications. We hope the over optimistic claims about agent technology are moderated. It is a sad fact that the only jarring voices in this hymn of confident approbation about agent technology come from those doing the *real* research, and who know what some of the real technical, social and ethical challenges are.

7 Conclusion and a Postscript

This paper has pilfered from a diverse literature in order to overview, in a single paper, the rapidly evolving area of software agents. Only one other paper (Wooldridge & Jennings, 1995a) has attempted a similar extensive review of this area, which they do from a theories, architecture and languages angle. In this paper, we have overviewed the same area from the viewpoint of the clear diversity of agents being investigated in universities and research laboratories worldwide. We hope it provides a useful contribution to understanding this exciting field of software agents. However, we conclude this paper with the following, arguably controversial and/or casual, postscript.

The word 'agent' is currently in vogue in the popular computing press as it is within the artificial intelligence (AI) and computer science communities. It has become a buzzword because it is both a technical concept and a metaphor. However, its rampant use should conjure up the problems faced with other flamboyant titles including 'artificial intelligence' itself: far too ambitious claims precede the real technical work that follows; the subject tends to attract dilettantes who just 'jump on the bandwagon' in order to make a "quick buck" or for some other ulterior motives; and coupled with another contentious word such as 'intelligent' as in 'intelligent agent', *prima facie* at least, it generates unrealistic expectations of the state-of-the-art. These exaggerated claims, over-expectations and the 'oversell' eventually drown out the real (and sometimes excellent) achievements, which never match the hype. In short, much is promised but little is delivered! The aftermath of all this is usually quite predictable: interest in the area soon wanes and the research funding, gradually or rapidly, disappears. Eventually, the dilettantes leave in search for 'greener pastures' and leave the real researchers to pick up the pieces. Is this the fate for agents research - a passing fad? Hopefully not!

There are several things which the serious agent researchers can do. Firstly, they can drop the 'intelligent' in intelligent agents as we have done in the title of this paper: its connotation, and hence expectations, are much less. Secondly, they could attempt to ensure, where possible, that dilettantes do not publish articles on agents in the popular press, at least not until it has been reviewed by someone whose interest in the area is more than superficial. This is sometimes possible because some experts usually get asked to review such articles before they go to press. Thirdly, they (i.e. experts) should not *themselves* engage in overselling the domain and, lastly, they must always be critical of the progress in the area in order to provide a more realistic appraisal of the state-of-the-art. In this overview paper, we have attempted to abide by the principles of this doctrine. It is up to the reader to judge how successful we have been in meeting these principles in this paper.

8 Acknowledgements

We acknowledge Nader Azarmi for encouraging us to write this paper, particularly through his provision of a fraction of the literature referenced in it. David Griffiths also encouraged the writing of this article. Several informal discussions with Barry Crabtree, Mark Wiegand, Paul O'Brien, Robin Smith and Nader Azarmi have also shaped some of the views propounded in it. Barry, Robin, Nader, Divine Ndumu and Brian Odgers have also provided feedback which has improved the quality of this paper. Lastly, we acknowledge the comments of the three anonymous reviewers, though two of them were not so anonymous afterall. The author bears all the responsibility for any misunderstandings and/or errors therein. The views expressed are also those of the author, and not necessarily those of BT plc. This work was funded by BT Laboratories.

... Last Section

Bibliography ...