

Agent Technology Roadmap: Overview and Consultation Report

agent based computing

AgentLink III, December 2004

Introduction

AgentLink III

AgentLink III is an Information Society Technologies (IST) Coordination Action for agent-based computing, funded under the European Commission's Sixth Framework Programme (FP6), running through 2004 and 2005. Agent-based systems are one of the most vibrant and important areas of research and development to have emerged in information technology in recent years, underpinning many aspects of broader information society technologies.

The long-term goal of AgentLink is to put Europe at the leading edge of international competitiveness in this increasingly important area. AgentLink is working towards this by seeking to achieve the following objectives.

- To gain competitive advantage for European industry by promoting and raising awareness of agent systems technology.
- To support standardisation of agent technologies and promote interoperability.
- To facilitate improvement in the quality, profile, and industrial relevance of European research in the area of agent-based computer systems, and draw in relevant prior work from related areas and disciplines.
- To support student integration into the agent community and to promote excellence in teaching in the area of agent-based systems.
- To provide a widely known, high-quality European forum in which current issues, problems, and solutions in the research, development and deployment of agent-based computer systems may be debated, discussed, and resolved.
- To identify areas of critical importance in agent technology for the broader IST community, and to focus work in agent systems and deployment in these areas.

Purpose of the Roadmap

In trying to raise awareness and to promote take-up of agent technology, there is a need to inform the various audiences of the current state-of-the-art and to postulate the likely future directions the technology and the field will take. This is needed if commercial organisations are to best target their investments in the technology and its deployment, and also for policy makers to identify and support areas of particular importance. More broadly, presenting a coherent vision of the development of the field, its application areas and likely barriers to adoption of the technology is important for all stakeholders. AgentLink is undertaking this technology roadmapping study in order to develop just such a strategy for agent research and development.

The roadmap is a living document, and the roadmapping process is ongoing. This interim report reflects the initial efforts of the roadmapping team in developing a framework and outlining key directions and messages that have been elicited to date. Feedback on all aspects of this document is welcome; send email to the roadmap coordinator, Michael Luck, at mml@ecs.soton.ac.uk.

What is agent technology?

Agents can be defined to be autonomous, problem-solving computational entities capable of effective operation in dynamic and open environments. Agents are often deployed in environments in which they interact, and sometimes cooperate, with other agents (including both people and software) that have possibly conflicting aims. Such environments are known as multi-agent systems. Agents can be distinguished from objects (in the sense of object-oriented software) in that they are autonomous entities capable of exercising choice over their actions and interactions, and may act to achieve individual objectives. They are able to exercise autonomy by choosing how to perform the tasks assigned to them or by deciding on operational tasks to satisfy user objectives. More importantly, they make these choices in the context of dynamic environments in which they are deployed. Agents cannot, therefore, be directly invoked like objects but can be assigned tasks by their owners. Nevertheless, they may be constructed using a wide range of technologies, including object technology, Web Services and others. These notions find application in relation to several distinct aspects, considered in turn below.

Agents as design metaphor

Agents provide designers and developers with a way of structuring an application around autonomous, communicative components, and lead to the construction of software tools and infrastructure to support the design metaphor. In this sense, they offer a new and often more appropriate route to the development of complex systems, especially in open and dynamic environments. In order to support this view of systems development, particular tools and techniques need to be introduced. For example, methodologies to guide analysis and design are required, agent architectures are needed for the design of individual components, tools and abstractions are required to enable developers to deal with the complexity of implemented systems, and supporting infrastructure (including more general, current technologies, such as Web Services) must be integrated.

Agents as a source of technologies

Agent technologies span a range of specific techniques and algorithms for dealing with interactions in dynamic, open environments. These include issues such as balancing reaction and deliberation in individual agent architectures, learning from and about other agents in the environment, eliciting and acting upon user preferences, finding ways to negotiate and cooperate with other agents, and developing appropriate means of forming and managing coalitions. Moreover, the adoption of agent-based approaches is increasingly influential in other domains. For example, multi-agent systems have already provided faster and more effective methods of resource allocation in complex environments, such as the management of utility networks, than any human-centred approach.

Agents as simulation

Multi-agent systems offer strong models for representing real-world environments with an appropriate degree of complexity and dynamism. For example, simulation of economies, societies and biological environments are typical application areas.

The use of agent systems to simulate real-world domains may provide answers to complex physical or social problems that would be otherwise unobtainable, as in the modeling of the impact of climate change on biological populations, or modeling the impact of public policy options on social or economic behaviour. Agent-based simulation spans: social structures and institutions to develop plausible explanations of observed phenomena, to help in the design of organisational structures, and to inform policy or managerial decisions; physical systems, including intelligent buildings, traffic systems and biological populations; and software systems of all types, currently including eCommerce and information agency.

Agent systems are not simply panaceas for these large problems; they have been demonstrated to provide concrete competitive advantages such as:

- improving operational robustness with intelligent failure recovery;
- reducing sourcing costs by computing the most beneficial acquisition policies in online markets; and
- improving efficiency of manufacturing processes in dynamic environments.

Roadmap Consultation Report

This consultation report aims to provide an overview of the areas and issues that will be covered in the full AgentLink roadmap, to be published in the autumn of 2005, and to act as a call to relevant communities to contribute to its development. It begins with a review of recent technological innovations that provide a context and infrastructure for the development of agent systems and motivate the need for the use of agent technologies. In addition, several trends and drivers have also recently emerged that suggest the use of agents, and are pushing forward development both in their own areas (in many cases through the use of agents) and in the field of agent-based computing. Timescales for agent technologies are then considered, through early survey results, the aggregation of relevant areas from Gartner's Hype Cycle, and prior analyses of the field, as well as the European position, through an assessment of activity in recent years. Then, the report enumerates several major research areas and identifies key challenges. Finally, the report establishes some initial recommendations for actions required to move agent-based computing to the next level of maturity and deployment.

Further information on the roadmap activity and on how you may be involved in AgentLink's roadmapping effort can be found at the end of this document.

Technological Context

Until recently, one of the major hurdles for the widespread adoption of agent technologies within the wider IT context has been the lack of infrastructure able to support the creation of dynamic and heterogeneous networks of devices and services that is central to the support of significant agent-based systems.

However, recent years have seen the development of a vast array of middleware technologies to support emerging enterprise level (and quality of) systems. Such technological infrastructure ranges from low-level wireless communication protocols like Bluetooth to higher-level Web Services technologies. In addition, they also span the range of devices supported from limited-capability devices such as mobile phones and PDAs to workstations and high-performance computing.

These positive moves with regards to the technological context for agents, are illustrated in Figure 1, where the years in which the main technologies that can facilitate agent-based systems development are noted. While research in agent technologies has been active for over a decade now, the figure shows that it is only as of 1999, with the appearance of effective service-oriented technologies (Jini) and pervasive computing technologies (Bluetooth) that truly dynamic networked systems could be built without big investments in establishing the underlying infrastructure. In particular, as of 2002, with the emergence of Grid computing and calls for adaptive wide-scale Web Service-based solutions there is now a real need to provide attractive solutions to the higher-level issues of communication, coordination and security.

Web services technologies have emerged out of a realisation that alternative approaches, such as CORBA or Java RMI, were too tightly coupled to enable truly open distributed systems. Web Services provide a platform and programming language independent solution that can deal with the exchange of messages, the description of services and their publication and discovery. Each of these issues is a necessary component for any agent-based architecture.

Other efforts can also play a significant role in support of agent-based systems. Some of the more relevant ones are described below.

- Jini is a set of network technologies and protocols aimed at facilitating the dynamic discovery of online services. It is a Java-centric technology that is based around the notion of federated lookup services, which provide access to services registered with the lookup services.
- UPnP is aimed at enabling the integration of devices based on a set of technologies that make minimal requirements about the underlying network infrastructure to support interaction.
- The JXTA project is an attempt to provide a set of protocols for peer-to-peer communication that are platform and application neutral using XML.
- Bluetooth is an open standard for wireless communication based on a low-

AGENTLINK ROADMAP

Agent Technology: Overview and Consultation Report

cost short-range radio link. It has a particularly significant role to play in facilitating the discovery of limited-capability devices and dynamic networking.

- Other technologies such as Gnutella for file sharing, or P2P file sharing, all provide potential infrastructure and services that can be enhanced through agent technologies.

In general, it is clear that current technological developments are increasingly addressing problems that have long been posed within the agent research community. In particular, they are responding to some of the underlying infrastructural needs for agent-based systems, such as a standardised means for discovery and communication between heterogeneous services. This suggests two clear and inter-related trends. First, supporting technologies are emerging very quickly. Thus, to some extent, the research problem has moved from one of infrastructure to one of the higher-level issues relating to effective coordination and cooperation between services. Second, a very large number of systems are being built and designed using these emerging infrastructures, and are becoming ever more agent-like — their developers face the same problems that the agent community has encountered in the past.




Figure 1: Agent-related technologies for infrastructure support

Emerging Trends and Critical Drivers

Trends

Web Services

Web Services provide a standard means of interoperating between different software applications, running on a variety of different platforms. They are software systems that aim to support interoperable machine interaction via networks, with interfaces described in a machine-processable format. Other systems interact with the Web Service in a manner prescribed by Web-related standards.

In addition to their interoperability and extensibility (arising from the use of XML), the power of Web Services is that they can be combined in a loosely coupled way in order to achieve complex operations. Programs providing simple services can interact with each other in order to deliver more sophisticated services. In this sense, Web Services must be realised by concrete agents that send and receive messages, while the services are the resources characterised by the functionality provided. In the same way as agents perform tasks on behalf of a user, Web Services provide this functionality on behalf of owners, people or organisations.

Web Services thus provide a ready-made infrastructure that is almost ideal for use in supporting agent interactions in a multi-agent system. More importantly, perhaps, it is widely accepted, standardised, and likely to be the dominant base technology over the coming years. Conversely, as Web Services have become more complex, many results from the agent community have begun to be relevant to the design of Web Services systems.

Grid Computing

The Grid is the high-performance computing infrastructure for supporting large-scale distributed scientific endeavour that has recently gained heightened and sustained interest from several communities [3]. It provides a means of developing eScience applications such as those demanded by, for example, the Large Hadron Collider facility at CERN, engineering design optimisation, bioinformatics and combinatorial chemistry. Yet it also provides a computing infrastructure for supporting more general applications that involve large-scale information handling, knowledge management and service provision. Typically, Grid computing is abstracted into several layers, which might include: a data-computation layer dealing with computational resource allocation, scheduling and execution; an information layer dealing with the representation, storage and access of information; and a knowledge layer, which deals with the way knowledge is acquired, retrieved, published and maintained.

The Grid thus refers to an infrastructure that enables the integrated, collaborative use of high-end computers, networks, databases, and scientific instruments owned and managed by multiple organisations. Grid applications often involve large amounts of data and computing and often require secure resource sharing across organisational boundaries; they are thus not easily handled by today's Internet and Web infrastructures.

Drivers

Ambient Intelligence

The notion of Ambient Intelligence has largely arisen through the efforts of the European Commission in identifying challenges for European research and development in Information Society Technologies. Aimed at seamless delivery of services and applications, it relies on the areas of ubiquitous computing, ubiquitous communication and intelligent user interfaces. The vision describes an environment of potentially thousands of embedded and mobile devices (or software artifacts) interacting to support user-centred goals and activity. This suggests a component-oriented view of the world in which the artifacts are independent and distributed. The consensus is that autonomy, distribution, adaptation, responsiveness, and so on, are key characterising features of the components, and in this sense they share the same characteristics as agents.

Ambient Intelligence requires these agents to be able to interact with numerous other agents in the environment around them in order to achieve their goals. Such interactions take place between pairs of agents (in one-to-one cooperation or competition), between groups (in reaching consensus decisions or acting as a team), and between agents and the infrastructure resources that comprise their environments (such as large-scale information repositories). Interactions like these enable the establishment of electronic institutions or virtual organisations, in which groups of agents come together to form coherent groups able to achieve overarching goals.

The environment provides the infrastructure that enables Ambient Intelligence scenarios to be realised. On the one hand, agents offering particular services can be distinguished from facilitating services such as the physical infrastructure needed to support effective interaction through sensors and actuators, and the physical connectivity for supporting quick and efficient interactions, for example. On the other, they can also be distinguished from the virtual infrastructure needed to support resource discovery, large-scale distributed and robust information repositories (as mentioned above), and the logical connectivity needed to enable effective interactions between large numbers of distributed agents and services, for example.

In relation to pervasiveness, it is important to note that scalability (more particularly, device scalability), or the need to ensure that large numbers of agents and services are accommodated, and heterogeneity of agents and services is facilitated by the provision of appropriate ontologies to enable effective interactions. Addressing all of these aspects will require efforts to provide solutions to issues of operation, integration and visualisation of distributed sensors, ad hoc services and network infrastructure.

Autonomic Computing

In response to the problems relating to the explosion of information and integration of technology into everyday life, and the associated problems of complexity in managing and operating computer systems, Autonomic Computing draws on the autonomic function of the human central nervous system, which controls key functions without conscious awareness or involvement. The paradigm first proposed by IBM [2], Autonomic Computing is an approach to self-managed computing systems with a minimum of human interference.

Its goal is a network of sophisticated computing components that give users what they need, when they need it, without a conscious mental or physical effort. Among the defining characteristics of an autonomic system are the following: it must automatically configure and reconfigure itself under varying (and unpredictable) conditions; it must seek to optimise its operation, monitoring its constituent parts and fine-tuning its workflow to achieve system goals; It must be able to discover problems and recover from routine and extraordinary events that might cause malfunctions; it must act in accordance with its current environment, adapting to best interact with neighbouring systems, by interacting and negotiating for resource use with other systems; it must function in a heterogeneous world and implement open standards. it must marshal resources to reduce the gap between its (user) goals and their achievement, without direct user intervention.

Ultimately, the goal is to realise the promise of IT: increasing productivity while minimising complexity for users. The key message to be drawn from this vision is that it shares many of the goals of agent-based computing, and agents offer a way to manage this complexity. Undoubtedly agent technologies will be required to deliver the goal of autonomic computing.

Semantic Web

Since it was first developed in the early 1990s, the World Wide Web has rapidly and dramatically become a critically important and powerful medium for communication, research and commerce. However, the Web was designed for use by humans, and its power is limited by the ability of humans to navigate the data of different information sources.

The Semantic Web is thus based on the idea that the data on the web can be defined and linked in such a way that it can be used by machines for the automatic processing and integration of data across different applications. This is motivated by the fundamental recognition that, in order for the Web to scale, programs must be able to share and process data, particularly when they have been designed independently. The key to achieving this is by augmenting web pages with descriptions of their content so that it is possible to reason about that content.

Among the particular requirements for the realisation of the Semantic web vision are: rich descriptions of media and content to improve search and management; rich descriptions of Web Services to enable and improve discovery and

“The Semantic Web is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation.” — Tim Berners-Lee, James Hendler, Ora Lassila, The Semantic Web, Scientific American, May 2001

composition; common interfaces to simplify integration of disparate systems; and a common language for the exchange of semantically-rich information between software agents.

It should be clear from this that the Semantic Web demands effort and involvement from the field of agent-based computing, and is intimately tied up with it. Indeed, it offers a rich breeding ground for both further fundamental research and a whole range of agent applications that can (and should) be built on top of it. In the first stage of Web development, all Web-pages were created by humans and read by humans. In the second (current) stage, most Web-pages are created by machines (for example, automatically generated from database queries) but still read by humans. In the next stage (the Semantic Web), most Web-pages will be created by machines and read by machines.

Complex Systems

Modern software and technological systems are among the most complex human artifacts, and are ever-increasing in complexity. Some of these systems – for example, the Internet – were not designed but simply grew organically, with no central human control or even understanding. Other systems, such as global mobile satellite communications networks or current PC operating systems, have been designed centrally, but comprise so many interacting components and so many types of interactions, that no one person or even team of people could hope to comprehend the detailed system operations. This lack of understanding may explain why such systems are prone to error as, for example, in the large-scale electricity network failures in North America and in Italy two years ago.

Moreover, many systems that affect our lives involve more than just software. For example, the ecosystem of malaria involves natural entities (parasites and mosquitos), humans, human culture, and technological artifacts (drugs and treatments), all interacting in complex, subtle and dynamic ways. Intervening in such an ecosystem – for example, providing a new treatment regime for malaria – may have unintended and unforeseen consequences due to the little-understood nature of these interactions. The science of complex adaptive systems is still in its infancy, and provides little yet in the way of guidance for designers and controllers of specific complex systems.

Whether such complex, adaptive systems are designed or not, their management and control is vitally important to modern societies. Agent technologies provide a way to conceptualise these systems as comprised of interacting autonomous entities, each acting, learning or evolving separately in response to interactions in their local environments. Such a conceptualisation provides the basis for realistic computer simulations of the operation and behaviour of the systems, and of design of control and intervention processes [1]. For systems that are centrally designed, such as electronic markets overlaid on the Internet, agent technologies also provide the basis for the design and implementation of the system itself. Indeed, it has been argued that agent technologies provide the only way to cope with the increasing complexity of modern software systems [8] – pervasive

devices, ambient intelligence, continuous operation (allowing no down-time for upgrades or maintenance), and open systems.

Summary

It is natural to view large systems in terms of the services they offer, and consequently in terms of the entities or agents providing or consuming services. The domains discussed here reflect the trends and drivers for applications in which typically many agents and services may be involved, and spread widely over a geographically distributed environment. Figure 2 depicts the emergence of these driver domains over time, illustrating the currency, and suggesting that their maturity, which will demand the use of agent technologies, is likely to be some years away.

Most importantly perhaps, the environments that have been identified here are open and dynamic so that new agents may join and existing ones leave. In this view, agents act on behalf of service owners, managing access to services, and ensuring that contracts are fulfilled. They also act on behalf of service consumers, locating services, agreeing contracts, and receiving and presenting results. In these domains, agents will be required to engage in interactions, to negotiate, and to make pro-active run-time decisions while responding to changing circumstances. In particular, agents will need to collaborate and to form coalitions of agents with different capabilities in support of new virtual organisations.

Of course, these drivers say nothing specific about a whole range of specific areas with the field of agent-based computing, including human-agent interfaces, learning agents, robotic agents, and many others, but they provide a context that is likely to drive forward the whole field.




Figure 2: The emergence of agent-related domains over time.

Timescales for Deployment of Agent Systems

Research and development of agent technologies still has a long way to go. We are already seeing the use of agent technologies in particular applications and technology propositions, but this is just the first wave of very early adopters in sample contexts. For example, Magenta Technologies is a UK-based company that specialises in the application of agents to optimising logistics; it is undertaking early trials for a major oil-shipping client, but the system is not yet fully deployed. In the US, companies such as Nutech have already deployed systems for clients that include Air Liquide and Southwest Airlines in similar contexts.

It is still too early to consider the penetration of different industry sectors, but the AgentLink survey of key technology experts identified the industry sectors of telecommunications and networks, manufacturing, and transport as the three most likely to have the greatest impact from the application of agent systems. Interestingly, as a domain for impact of agents, computer software came relatively low down the list, after healthcare and on a par with wholesale and retail trade. This contrasts with much academic research which has been focussed on eCommerce and eBusiness systems in recent years, partly because of its relative currency in the light of the Internet boom, and partly because of its ready availability as a domain to study. One question to consider is whether this focus is out of sync with business application realities, or whether the survey points beyond immediate application domains to those where there is likely to a fundamental change in modality.

The early results of the survey also reveal that it is likely to be another six years before we see more general deployment. The mean expectation for mainstream deployment of general agent applications was the year 2011, and limited but identifiable deployment of agent technologies (such as negotiation) as part of eCommerce applications) being achieved by 2007. Similarly, the respondents to the survey do not expect development costs in agent technologies to be offset by revenues generated before 2009.

One might suggest that even this is optimistic, especially the latter prediction. At the same time as making these predictions, the survey reveals that in terms of academic research, more is needed. The vast majority of experts believe that a medium term vision and commitment of 4–6 years was sensible for development of relevant core techniques, with a substantial number of respondents biasing towards a longer-term timescale of 7–10 years.

These snippets of information are not yet overly illuminating, but more complete results of the survey will be released in 2005, and the survey will be opened to more general contribution, with complete results being published in the full roadmap.

¹ The first round of a Delphi survey aimed at supporting the AgentLink Roadmapping effort, carried out in late 2004.

Activity Landscape and European Position

Activity in Europe

The European position on research and development in agent systems is healthy. There have been numerous active research groups in universities and research laboratories across Europe since the early days of the emergence of the field of agent-based computing as a distinct discipline, and the quality of work done is competitive at a global level. One reason for this is that since 1998, the European Commission has provided funding (albeit limited) to support the community through coordination projects, providing a focus and coherence to the community that might not otherwise have been possible. The value of these AgentLink projects has not just been in academia; AgentLink II, which ran from 2000 to 2003, counted around 40% of its organisational members from industry. Interestingly, research activity was generally sustained despite the trough caused by the bursting of the Internet bubble, and it can be argued that the efforts of the Commission in supporting the agent community helped to minimise its consequences.

Yet, there have been consequences. According to [15], in the period before the bursting of the bubble, the ICT sector was characterised by hypercompetition, in which industries tried to outpace their competitors with speed of innovation. Business innovations were implemented in a quick and dirty fashion so as to minimise time to market and achieve rapid, exponential growth, at the cost of poorly conceived business models, and a high cash burn rate. The collapse led to consolidation in ICT sectors, and the emphasis has since shifted to the e-enablement of core business processes, like fully integrated supply chains and supply networks, with a focus on visible and measurable impact. This shift can now also be seen in the positioning of agent technology providers, who now focus more on these latter areas, and less on fundamental process change.

The European Position

In the USA, ICT is stimulated by the cultivation of a high-tech entrepreneurial culture, providing ready customers for new technologies and close cooperation between industry and universities. In addition, public R&D is oriented towards areas that are considered the application areas of the future and earmarked as national priorities. Among the USA's 16 "grand challenges" are the following relevant to agent technologies: knowledge environments for science and engineering; collaborative intelligence: integrating humans with intelligent technologies; and managing knowledge intensive organisations in dynamic environments.

By contrast, European innovation culture and policy are more sluggish, despite the efforts of the European Commission. The grand challenges may be reflected in the strategic objectives of FP6, and in other relevant policy documents, but the ready customers for new technologies and the close cooperation between business and universities are not always apparent. In addition, there is also a recognition at the level of the European presidency, in the report published by the Dutch Ministry of Economic Affairs [15], of the need to "accelerate the introduction of disruptive technologies," the most relevant of the 10 breakthroughs identified as being needed to move towards the

Lisbon goals [16]. Broad deployment and use of disruptive technologies require understanding and acceptance. Yet the lack of a proper dialogue between industry, government and society stakeholders often obstructs the process of achieving it.

However, due to the support of AgentLink by the Commission, at least some form of drawing together of the research and business communities has taken place in the domain of agent-based computing, and there are ready channels for interaction to facilitate rather different models of cooperation.

Figure 5 illustrates the impact of activity in Europe, with AgentLink and Agentcities.NET providing coordination of the community through a period of intense change and innovation at the research level. Usable FIPA standards, for example, were developed in 1998, but matured in 2000; several FIPA compliant agent platforms (JADE, Zeus and FIPA-OS) were also released by 2000, while other systems have also made an impact, such as the Creatures computer game, for example. Meanwhile developments in the Semantic Web gave rise to OIL and then DAML+OIL. At the bottom of the figure, key events in the development of the research community are indicated: the International Conference on Multi-Agent Systems appeared in 1995, the Autonomous Agents Conference in 1997, and both were combined into the International Joint Conference on Autonomous Agents and Multi-Agent Systems in 2002. In addition, the International Foundation for Multi-Agent Systems was set up in 1998, and a similar European initiative was launched in 2003 with a European workshop.




Figure 5: European activity in agent-based computing in recent years

The Agent Technology Hype Cycle

Technology forecasting is a notoriously difficult task. In seeking to understand patterns of technology development in the mid-1990s, Gartner devised a model known as the Hype Cycle (described below), which indicates the maturity of a technology, from initial excitement to disillusionment and then, for some eventual market acceptance. Gartner's July 2004 analysis of technologies and applications [7–12] places various agent technologies, agent-related technologies, application domains and drivers at various different points in the cycle.

In terms of infrastructure, business process execution languages (BPEL) are rising on the technology trigger path, with between 1% and 5% market penetration. Basic Web Services for service definition and application integration, using SOAP and WSDL, are climbing the slope of enlightenment and are implemented by major software vendors, reaching 20% to 50% market penetration. Advanced Web Services for higher quality of service, which will enable advanced business-critical functions over standards-based networks, using SOAP, WSDL, UDDI, WS-Security and WS-R, depend on the availability of standards, and implementations are not yet fully delivered by vendors.

Drivers and domains figure primarily through the Semantic Web, which is placed at the peak of expectation; while the expectation is for a transformational impact, at present it has less than 1% market penetration. Similarly, the Trading Grid, an interconnection of networks and marketplaces to support virtual organisations, is also transformational but just at the very start of the cycle. With lower perceived impact, but more mature, are eMarketplaces, now with up to 5% market penetration. Each of these is predicted to take up to 10 years to plateau.

Intelligent agents as a whole are seen as being in the trough, having been over-hyped in the past, like synthetic characters, or chatterbots. By contrast, web self-service agents, which act on a customer's or business's behalf to automate transactions are finally "catching on", and have reached up to 5% penetration. In all these cases, however, these are lightweight agents, with the mainstream of agent technologies still to engage. For example, agent-based integration is concerned with enabling distributed applications that demand autonomy and flexibility. In this area, commercial technology is still new, and the sector is dominated by small startups and only a small number of users, so agent-based integration is at the start of the cycle. Gartner estimates that market penetration is less than 1% of the target. Given the position of the Semantic Web, this is perhaps not surprising, but the time to plateau is shorter, at up to 5 years.

At the embryonic stage are: swarm intelligence, or emergent computing, which fits directly with the complex systems discussed above; and affective computing, which seeks to recognise human emotional states for better user interfaces. At present, these are mainly in the domain of research laboratories

The Hype Cycle for agent technologies, based on Gartner's analysis across different domains is illustrated in Figure 3.

The Hype Cycle involves the following five stages.

1. Technology trigger: introduction of the technology to a wider audience.
2. Peak of inflated expectations: the high point, at which the claims of the benefits of the technology are often exaggerated.
3. Trough of disillusionment: as the promises fail to be delivered, many observers begin to ignore the technology.
4. Slope of enlightenment: more is learned about the technology and, as many of the problems from the trough are resolved, standardisation takes place, and the technology is adopted primarily in the areas that perceive the greatest benefit.
5. Plateau of productivity: the new technology is well understood and stable, and becomes mainstream. Benefits and drawbacks for adoption are also widely known.




Figure 3: The agent technology Hype Cycle

Technology Roadmap

In any high-technology domain, the systems deployed in commercial or industrial applications tend to embody research findings somewhat behind the leading edge of academic research. Multi-agent systems are no exception to this, with currently-deployed systems having features found in published academic research and prototypes of three to five years ago. By looking at current academic research interests and areas of focus, we are able to extrapolate future trends in deployed systems.

Accordingly, we have identified four broad phases of the future development of deployed multi-agent systems. These phases are, of necessity, only indicative, since some companies and organisations will be leading users of agent technologies, pushing applications ahead of these phases, while many others will be laggards. We aim to describe the majority of research challenges at each time period. Note that this view on timescales takes the research view rather than the development view in that typically research is about three to five years ahead of development in this context. This analysis is an updated version of the analysis initially undertaken in [3].

Phase 1: Current

Multi-agent systems are currently typically designed by one design team for one corporate environment, with participating agents sharing common high-level goals in a single domain. These systems may be characterised as closed. The communications languages and interaction protocols are in-house protocols, defined by the design team prior to any agent interactions. Systems are usually only scalable under controlled, or simulated, conditions. Design approaches tend to be ad hoc, inspired by the agent paradigm rather than using any specific methodologies. Although this is still largely true, there is now an increased focus on, for example, taking methodologies out of the laboratory and into development environments, with commercial work being done on establishing industrial-strength development techniques and notations.

It remains true that, for the foreseeable future, there will be a substantial commercial demand for closed multi-agent systems because of the security concerns that arise from open systems. While progress in this respect will change the nature of agent systems, the importance of closed, well-protected systems must not be underestimated.

Phase 2: Near-Term Future

In the next phase of development, systems will increasingly be designed to cross corporate boundaries, so that the participating agents have fewer goals in common, although their interactions will still concern a common domain, and the agents will be designed by the same team, and will share common domain knowledge. Increasingly, standard agent communications languages, such as FIPA ACL, will

be used, but interaction protocols will remain non-standard. These systems will be able to handle large numbers of agents in pre-determined environments, such as those of Grid applications. Development methodologies will have reached a degree of maturity, and systems will be designed on top of standard infrastructures such as Web Services, for example. Example systems developed in this phase include those to enable automated scheduling coordination between different departments of the same company, closed user groups of manufacturing suppliers engaged in electronic procurement activities, or network-centric operations.

Phase 3: Medium-Term Future

In the third phase, multi-agent systems will permit participation by heterogeneous agents, designed by different designers or teams. Any agent will be able to participate in these systems, provided their (observable) behaviour conforms to publicly-stated requirements and standards. However, these open systems will typically be specific to particular application domains, such as B2B eCommerce or bioinformatics. The languages and protocols used in these systems will be agreed and standardised, perhaps drawn from public libraries of alternative protocols that will, nevertheless, likely differ by domain. In particular, it will be important to master this semantic heterogeneity.

	Short Term	Medium Term	Long Term
Industrial Strength software	Peer to peer Better development tools Web Services Agent UML	Generic Designs for Coordination Libraries for agent-oriented development	Best practice in agent systems design
Agreed Standards	FIPA ACL Peer to peer Better development tools Web Services Semantic description	Flexible business/trading languages Libraries of interaction protocols	Tools for evolutions of communications languages and protocols
Infrastructure for Open Communities	Web mining Semantic interaction Data integration and Semantic Web	Agent-enabled semantic web (services) Electronic institutions	Shared, improved ontologies Dynamic norms, roles, laws
Reasoning in Open Environments	Organisational views of agent systems	Enhanced understanding of agent society dynamics Theory and practice of argumentation strategies Norms and social structure Theory and practice of negotiation strategies	Automated eScience systems and other application domains
Learning Technologies	Adaptation Personalisation Hybrid technologies	Evolving Agents Self organisation Distributed learning	Run-time reconfiguration and re-design
Trust and Reputation	Security and verifiability for agents Reliability testing for agents Self-enforcing protocols	Norms and social structures Reputation mechanisms Formal methods for open agent systems Electronic contracts	Trust techniques for coping with malicious agents

Figure 4: Agent technology comprises numerous problem areas that will be addressed over different timescales

Systems will scale to large numbers of participants, although typically only within the domains concerned, but with particular techniques (such as domain-bridging agents), to translate between separate domains. Thus, for example, a multi-agent system for automated meta-analysis of research results in some area of biology will be able to utilise bridge agents to undertake commercial negotiations when interaction with an eCommerce system is required, say for access to information protected by patent. System development will proceed by standard agent-specific methodologies, including templates and patterns for different types of agents and organisations. Semantic issues related to, for example, coordination between heterogeneous agents and access control, are of particular importance here.

Examples of systems in this phase will be corporate B2B electronic procurement systems permitting participation by any supplier, rather than closed user groups.

Phase 4: Long-term Future

The fourth phase in this projected future will see the development of open multi-agent systems spanning multiple application domains, and involving heterogeneous participants developed by diverse design teams. Agents seeking to participate in these systems will be able to learn the appropriate behaviour for participation in the course of doing so, rather than having to prove adherence before entry. Although standard communications languages and interaction protocols will have been available for some time, systems in this phase will enable these to emerge by evolutionary means from actual participant interactions, rather than being imposed. Of course, such languages, protocols and behaviours may be mere refinements of previously-developed standards, but will be tailored to their particular contexts of use.

By this phase, systems will be fully scalable in the sense that they will not be restricted to arbitrary limits (on agents, users, complexity, etc). As with the previous phase, systems development will proceed by use of rigorous agent-specific design methodologies.

Technologies and Timescales

Arising from this picture of the future of agent research, we see a number of broad technological areas of research and development over the next decade. These are summarised in Figure 4, with each area being broken down into distinct technological issues that are organised into the short term, medium-term and long-term future, according to the points in time at which they will attract successful attention from the research and development communities. In particular, the tables suggest that long-term issues are worthy of strategic investment and effort while short-term issues are largely already addressed or are being addressed. Below, we consider each in turn; a much more detailed treatment of many of these issues can be found in [3, 4].

Industrial Strength Software

One of the most fundamental obstacles to the take-up of agent technology is the lack of mature software development methodologies for agent-based systems. Clearly, basic principles of software and knowledge engineering need to be applied to the development and deployment of multi-agent systems, but they also need to be augmented to suit the differing demands of this new paradigm. In particular, agent software engineering methodologies need to draw on insights gained from the design of economic systems, social systems, and complex engineering control systems.

Agreed Standards for Open Systems Development

Much of the agent standardisation effort has fallen to the Foundation for Intelligent Physical Agents (FIPA) and the Object Management Group (OMG). Importantly, as technologies converge, other non-agent standards are becoming increasingly relevant; over the next few years, there will be a much larger role for the more widely-adopted, Web Services standards, as indicated above.

Infrastructure for Open Agent Communities

At present, information agents exist in academic and commercial laboratories, but are not widely available in real world applications. The move out of the laboratory is likely to happen over the next ten years, but a much higher degree of automation than is currently available in dealing with knowledge management is needed for information agents. In particular, this demands new web standards that enable structural and semantic description of information; and services that make use of these semantic representations for information access at a higher level. The creation of common ontologies, thesauri or knowledge bases play a central role here, and merits further work on the formal descriptions of information and, potentially, a reference architecture to support the higher level services mentioned above.

Reasoning in Open Environments

Reasoning is a critical faculty of agents, but the extent to which it is needed is determined by the context. While reasoning in general is important, in open environments there are some specific concerns relating to heterogeneity of agents, trust and accountability, failure handling and recovery, and societal change. Work must be continued on the representation of computational concepts for the norms, legislation, authorities, enforcement, etc., that can underpin the development and deployment of dynamic electronic institutions. Similarly, current work on coalition formation for virtual organisations is limited, with such organisations largely static. The automation of coalition formation may be more effective at finding better coalitions than humans in complex settings, and is required for example, for Grid applications.

One enabler for this is negotiation, yet research into negotiation mechanisms that are more complex than auctions and game-theoretic mechanisms is still in its infancy. Research into argumentation mechanisms, for example, and the strategies appropriate for participants under them, is also needed before these mechanisms will achieve widespread deployment. In addition, many electronic institutions will be required to make decisions collectively, aggregating in some fashion the individual preferences or decisions of the participants. Research on the application to agent societies of social choice theory from political science and sociology is in its infancy, and considerably more work is needed here.

Agent Adaptation and Learning

Even though learning technology is clearly crucial for open and scalable multi-agent systems, it is still in early development. While there has been progress in many areas, such as evolutionary approaches and reinforcement learning, these have still not made the transition to real-world applications. Reasons for this can be found in problems of scalability and in user trust in self-adapting software. In the longer term, learning techniques are likely to become a central part of agent systems, while the shorter term offers application opportunities in areas such as interactive entertainment, which are not safety-critical.

Trust and Reputation

Although currently deployed agent applications often provide good security, for agents autonomously acting on behalf of their owner, several additional factors need to be addressed. First, considerable effort must still be put into issues of security in open agent systems. Efforts by other communities are tackling some aspects here, but more on specific agent security concerns needs to be done. Second, collaboration of any kind, especially in situations in which computers act on behalf of users or organisations, will only succeed if there is trust. To ensure this trust requires, for example, the use of: reputation mechanisms to assess prior behaviour; norms (social rules) and the enforcement of sanctions; and electronic contracts to represent and enforce agreements.

Beyond the technical challenges are issues relating to user trust in adopting agent technology in the first place.

Challenges

As IBM has stated in relation to Autonomic Computing, the vision of agent-based computing itself is enough to constitute a Grand Challenge, because of the need to bring together multiple technical and scientific disciplines as well as stakeholders across different sectors. The specific technical challenges continue to change as the field of agent-based computing advances and matures, and as related areas (like those discussed above) emerge and galvanise efforts that contribute to the general area. Inevitably, standards will continue to be critical, but it is not clear whether these should come from within the agent community or should emerge from more general computing infrastructure progress. (Recent relevant standards efforts are depicted in Figure 6.) Nevertheless, some key challenges have already been articulated in relevant areas. This section summarises some of the more salient challenges relevant to agent-based computing, drawing on the work of the AgentLink Technical Forum Groups and on [1, 2, 3, 4].




Figure 6: Standards activity in the area of agent-based computing

Specific Challenges

Service architecture

There is a need for an integrated service architecture providing a robust foundation for autonomous behaviour, in order to support dynamic services, and important negotiation, monitoring, and management patterns. This will aid the convergence of agent and Grid technologies

Trust negotiation and management

Sophisticated distributed systems are likely to involve action in the absence of strong existing trust relationships. While middleware addresses secure authentication, and there exist techniques for verification and validation, these do not consider the harder problems of estab-

lishing, monitoring, and managing trust in a dynamic, open system. As discussed above, we need new techniques for expressing and reasoning about trust, to enable interaction in dynamic and open environments.

Provenance

Today's distributed environments (including Grid, Web Services and agent-based systems) suffer from a lack of mechanisms to trace results and a lack of infrastructures to build up trusted networks. Provenance enables users to trace how a particular result has been arrived at by identifying the individual and aggregated services that produced a particular output. From both an academic and an industrial perspective, the research question is to design, formalise and implement an open provenance architecture. Such a provenance architecture should be scalable and secure; it must be open and promote interoperability.

Resource allocation and coordination

The coordinated, autonomic management of distributed resources requires new abstractions, mechanisms and standards in the face of multiple, perhaps competing, objectives from different parties. Most R&D effort to date has focused on allocation and coordination mechanisms drawn from human societies (for example, common auction protocols), but the processing power and memory advantages of computational devices mean that completely new mechanisms and protocols may be appropriate for automated interactions.

Negotiation

To date, work on negotiation has provided point solutions. There is a need for a solid theoretical foundation for negotiation that covers algorithms and negotiation protocols, while determining which bidding or negotiation algorithms are most effective. From the system perspective, behaviour arising through the interplay of different negotiation algorithms must be analysed, and determining what kind of negotiation to consider, and when, must be established. Finally, effective negotiation strategies and protocols that establish the rules of negotiation, as

well as languages for expressing service agreements, and mechanisms for negotiating, enforcing, and reasoning about agreements are also needed.

Learning and optimisation theory

While learning and adaptation has a long tradition of research, particular contexts raise new issues. In sophisticated autonomic systems, agents continually adapt to the environment of other agents, so that agents must adapt to each other, violating the assumptions of single-agent learning theories, and potentially leading to instabilities. Here, optimisation that assumes a stationary environment also fails pathologically, and new methods must be developed.

Virtual organisation formation and management

Virtual Organisations (VOs) have been identified as the means to release the power of the Grid, but well-defined procedures for determining when to form new VOs, how to manage them, and ultimately how and when to disband them, are still missing.

Service composition

While Web Service technologies define conventions for describing service interfaces and workflows, we need more powerful techniques for dynamically describing, discovering, composing, monitoring, managing, and adapting multiple services in support of virtual organisations, for example.

Semantic integration

In open systems, different entities will have distinct information models, demanding that techniques are developed for bridging the semantic gap. Advances are required in such areas as ontology definition, schema mediation, and semantic mediation.

Action Plan and Recommendations

AgentLink's role is to support and promote research, development and deployment of agent technologies across industry and academia. As a result of the various consultations that have led to the development of this technology roadmap, some clear messages can be extracted relating to the perceived priorities of different sectors in order to move agent-based computing to the next level to encourage its further exploitation.

- Establish networks of practitioners and researchers with common interests, from academia, R&D labs, and industry and commerce, and organise meetings on specific topics to enable sharing of knowledge and experience.
- Stimulate new research on problems of critical importance to the deployment of agent systems in the emerging computing environments.
- Commission training and development courses aimed at students and developers, to increase the supply of knowledgeable experts available to European industry and commerce.
- Influence national funding agencies to allocate funding aimed at both addressing core research challenges and encouraging uptake of agent technologies, in addition to trans-European support.
- Encourage early adopters of agent technology, especially ones with some risk.
- Develop catalogue of early adopter case studies, both successful and unsuccessful, providing analyses of reasons for success and failure cases, and identifying best practice for agent-oriented development and deployment.
- Support and disseminate efforts in standardisation, across agent-specific organisations and other agent-relevant organisations.

What AgentLink is doing

AgentLink is developing a more substantial technology roadmap [7] for agent-based computing focused on R&D across business sectors as well as through fundamental blue-sky research. This roadmap is an ongoing process through 2005, which will deliver a final report for use by IT industries, user organisations, research institutes, policy-makers and funding agencies. It aims to provide clear messages on action priorities for all these groups in order to help to coalesce the divergent efforts and activities.

The first stage of consulting key domain experts, primarily in industry, in identifying trends and directions has already been undertaken, and this report provides the initial results of that effort as a means of moving to the second stage that will use the initial findings to focus more detailed consultation across a broader range of opinion and insight.

AgentLink is already undertaking to try to address some of the recommendations above; see www.agentlink.org for more information.

What you can do

This document is only the first stage in the process of developing AgentLink's 2005 roadmap. Its purpose is twofold: first as a marker, laying down the preliminary findings of the initial work undertaken on the roadmap; and second as a call to the wider community to engage in the process of developing the content of the final, more substantial document. In particular, we call on our colleagues to contribute to the activity through comment and constructive criticism of these highlights, to identify omissions, key strategic directions and challenges, and relevant domains that may impact the development of agent technologies, may be influenced by them, or may provide a fertile breeding ground for the development of applications.

We welcome contact and contributions from the academic community, but especially from industry and commerce.

There is also scope for contributing to our Delphi survey, by completing the online form at <http://www.agentlink.org/survey/>.

The 2005-06 Roadmap will be a public document distributed to the academic research community, commercial R&D laboratories, industry and commerce, government, research councils, policy-makers, etc. Nearly 4000 copies of the 2003 Roadmap were distributed in hard copy, with thousands more being distributed electronically. Its impact was broad and contributed to policy at both European and national levels.

Timescale:

Public consultation and content solicitation period: January – May 2005

Roadmap public draft: July 2005

Final consultation period: July – August 2005

Roadmap formal publication: October 2005

References

- [1] S. Bullock and D. Cliff, Complexity and Emergent Behaviour in ICT Systems, Foresight Report, DTI, UK, 2004
- [2] I. Foster, N. R. Jennings and C. Kesselman, "Brain meets brawn: Why Grid and agents need each other" Proceedings of the Third International Conference on Autonomous Agents and Multi-Agent Systems, New York, USA, 8–15, 2004.
- [3] I. Foster and C. Kesselman (eds.), The Grid 2: Blueprint for a New Computing Infrastructure, Morgan Kaufmann, 2004.
- [4] J. O. Kephart and D. M. Chess, The Vision of Autonomic Computing, IEEE Computer, 36(1), 41–50, 2003.
- [5] M. Luck, P. McBurney and C. Preist, Agent Technology: Enabling Next Generation Computing (A Roadmap for Agent Based Computing), AgentLink, 2003.
- [6] M. Luck, P. McBurney and C. Preist, A Manifesto for Agent Technology: Towards Next Generation Computing, Journal of Autonomous Agents and Multi-Agent Systems, 9(3), 203–252, 2004.
- [7] R. Phaal, C. Farrukh and D. Probert, Technology roadmapping—a planning framework for evolution and revolution. Technological Forecasting and Social Change, 71, 5–26, 2004
- [8] F. Zambonelli and H. V. Parunak, "Signs of a revolution in computer science and software engineering," Engineering Societies for the Agents World (ESAW 2002). Lecture Notes in Artificial Intelligence 2577, 13–28, Springer, Berlin, 2003.
- [9] Hype Cycle for Application Integration and Platform Middleware, Gartner, 2004.
- [10] Hype Cycle for Application Development, Gartner, 2004.
- [11] Hype Cycle for Human-Computer Interaction, Gartner, 2004.
- [12] Hype Cycle for B2B CRM Technologies, Gartner, 2004.
- [13] Hype Cycle for the Knowledge Workplace, Gartner, 2004.
- [14] Hype Cycle for Supply Chain Management, Gartner, 2004.
- [15] Rethinking the European ICT Agenda: Ten ICT-breakthroughs for reaching Lisbon goals, Ministry for Economic Affairs, Directorate-General for Telecommunications and Post, The Netherlands, 2004.
- [16] European Commission – (2000): The Lisbon European Council -An agenda of economic and social renewal for Europe. Contribution of the European commission to the Special European Council in Lisbon, 23-24th March, DOC/00/7.

Acknowledgements and information sources

AgentLink would like to thank all the organisations and individuals who have contributed, directly or indirectly in providing content and opinion in the development of this document and of the activities that will take place in the future: University of Southampton; University of Liverpool; European Commission; Institut d'Investigació en Intel·ligència Artificial, CSIC; Universitat Politècnica de Catalunya; International Conference on Autonomous Agents and Multi-Agent Systems, 2004; European Conference on Artificial Intelligence, 2004; European Workshop on Multi-Agent Systems, 2004; IEEE Systems, Man and Cybernetics (SMC) Conference, 2004.

AgentLink Technical Forum Groups

Agent-Oriented Software Engineering; Agents in Bioinformatics; Agents Applied in Healthcare; Intelligent Information Agents for Web Economies; Law and Electronic Agents; Networked Agents; Programming Multi-Agent Systems; Self-Organisation in Multi-Agent Systems; Trust for Open Collaborative Agent Business Environments.

AgentLink Roadmap Development

Core Team

Prof Michael Luck, University of Southampton, UK
Dr Peter McBurney, University of Liverpool, UK
Dr Onn Shehory, IBM Israel
Dr Steve Willmott, Universitat Politècnica de Catalunya, Spain

Management Committee

Monique Calisti
Wiebe van der Hoek
Michael Luck
Peter McBurney
Jorg Muller
Andrea Omicini
Terry Payne
Michal Pechoucek
Onn Shehory
Simon Thompson
Steve Willmott

Other Contributors

Ronald Ashri
Ulises Cortes
Mark d'Inverno
Jez McKean
Luc Moreau
Steve Munroe
Pablo Noriega
Juan A. Rodriguez
Hayden Shorter
Carles Sierra

AgentLink Support

Catherine Atherton
Becky Earl
Adele Maggs
Serena Raffin

