

## EXPRESSION OF INTEREST

### AgentLink III: A Research Network of Excellence for Agent-Based Computing

<b>Contact Details</b>	Dr Michael Luck Department of Electronics and Computer Science University of Southampton Southampton SO17 1BJ United Kingdom Email: <a href="mailto:mml@ecs.soton.ac.uk">mml@ecs.soton.ac.uk</a> Telephone: +44 23 8059 6657 Fax: +44 23 8059 3313
<b>Instrument</b>	Network of Excellence
<b>Scale</b>	20 Core Partners, 100 Associates, 50 Others; 4-5 years; Euro16-20 million
<b>Impact</b>	A network of excellence spanning all aspects of agent-based computing, in support of novel computing infrastructure and applications, including the Semantic Web, Grid computing and Ambient Intelligence.

#### 1 Objectives

Agent-based computing is one of the most vibrant and important areas of research and development to have emerged in the last 15 years. An agent is a computer system capable of flexible autonomous action in dynamic, unpredictable, typically multi-agent domains. While agent technology has begun to mature and has already been applied in the development of numerous systems, the newly emerging information infrastructure that will involve electronic and mobile commerce, open Grid computing, the Semantic Web, pervasive computing and other innovations holds the promise of a realisation of the Ambient Intelligence (AmI) vision only if supporting abstractions, technologies and tools can be developed. Common to all these areas are the underlying notions of open agent systems, which provide the entities involved either as enablers for the infrastructure or to construct applications on top of the infrastructure. In many cases, agents will be needed at several levels.

Many observers believe that agents represent the most important new paradigm for software for the next 10 years. Already, the concept of an intelligent agent has found currency in a diverse range of sub-disciplines of information technology, including computer networks, software engineering, object-oriented programming, artificial intelligence, human-computer interaction, distributed and concurrent systems, mobile systems, telematics, computer-supported cooperative work, control systems, and electronic commerce. Yet the challenges raised by the newer technological environment requires further development of both fundamental agent abstractions, algorithms, tools and techniques.

The agent metaphor is pervasive and touches many diverse fields, problems, and sources of research knowledge. It is also becoming increasingly important for understanding the complex software systems that are now being built. The agent paradigm (or *metaphor*) enables abstraction for

- users (to facilitate and ease *use* of advanced systems)
- complex software programs (for design and management at the right abstraction level)
- open network environments (of complex, active components – Grids, web services, etc.)

#### **Agent-Oriented System Design and Management: *Developing industrial strength agent software and tools***

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. Basic principles of software and knowledge engineering need to be augmented to suit the differing demands of agent systems. Most existing agent applications are developed in an *ad hoc* fashion, with limited specification of the requirements or design. To develop methods with which both the requirements of such systems, and the systems themselves, can be modelled and specified at a conceptually acceptable level of detail, characteristics of real-world multi-agent applications need to be identified, in relation to specific domains. At a tool level, software developers require sophisticated yet easy-to-use agent-oriented CASE environments to help in all aspects of system development.

Initial efforts have attempted to develop e.g., Agent UML with some success, but there now needs to be support for developers through industrial strength tools and community building activities. The challenges are technological in terms of tool support, methodological in providing ways to use the tools to support overarching development of agent systems, and societal in raising awareness and providing training support through, for example, a stock of case-studies that is resonant with developers. Importantly, the success of future developments is likely to be ensured not by considering agents in isolation, but through their integration with evolving (and current) system integration technologies (such as Jini, eSPEAK and UDDI). Agent technologies are particularly relevant at higher levels of interaction relating to communication, ontologies, content and semantics, whereas business integration frameworks focus on the provision of scalable and robust solutions to the lower levels, including protocols, syntax, distributed computing APIs, directory services etc. It is important to build on current efforts at ensuring that these are interoperable.

#### **Communication, Semantics and Open Systems: *Enabling high level interactions between tomorrow's business systems***

FIPA has defined a series of comprehensive agent standards covering key areas from message

transport (over multiple network infrastructures) and management services (directories, service naming) to agent communication (communication semantics, protocols and content languages). While its Agent Communication Language (ACL) is fast becoming a *de facto* standard, there is still a need for large-scale FIPA-compliant systems and demonstrators. Also required for open systems development will be libraries of interaction protocols designed for specific interactions, as more sophisticated interactions become common. These may use existing agent communication languages, as do various auction protocols when implemented using FIPA ACL, but may be implemented in *ad hoc* communications languages, as many dialogue game protocols for agent argumentation currently tend to do. In addition, open agent societies will require the ability to collectively evolve languages and protocols specific to the application domain and to the agents involved. Some work has commenced on defining minimum requirements for a group of agents with no prior experience of each other to evolve a sophisticated communications language, but this work is still in its infancy. Research in this area will draw on linguistics, social anthropology, biology, the philosophy of language and information theory.

Moreover, a much higher degree of automation than is currently available in dealing with knowledge management is needed. 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 is central, and merits further work on formal descriptions of information and, potentially, a reference architecture to support higher level services as above.

**Towards Smart Agent Systems: Developing infrastructure and reasoning capabilities for tomorrow's smart applications** *Organisational* approaches do not adequately handle the issues inherent in *open* multi-agent systems, namely heterogeneity of agents, trust and accountability, failure handling and recovery, and societal change. Human societies have successfully coped with similar issues, by creating institutions that establish norms for group dynamics in open systems. Agent-based computing needs to develop appropriate representations of analogous computational concepts to the norms, legislation, authorities, enforcement, etc., that can underpin the development and deployment of dynamic electronic institutions. Virtual organisations involve dynamic coalitions of small groups that can provide more services and make more profits than an individual group. Moreover, such coalitions can disband when they are no longer effective. At present coalition formation for virtual organisations is limited, with such organisations largely static. The automation of coalition formation will save both time and labour, and may be more effective at finding better coalitions than humans in complex settings. Although addressed in game theory for some time, it has typically been centralised and computationally infeasible, only applicable for small numbers of agents, and generally favouring one big coalition, limiting application scope. Emerging computation infrastructures such as the Grid now provide a greater need for effective work in virtual organisations to facilitate higher-level applications.

To date, related research into negotiation can be considered *point work*, with particular efforts or examples rather than a more coherent science of negotiation strategy. Strategies identified by economic or game theoretic reasoning, for example, tend to be specific to the auction or game mechanism involved. This makes their identification and deployment difficult, without any over-arching and computational theory. Consequently, it is not yet possible to define a computational agent capable of effective negotiation in any context. Moreover, research into complex negotiation and deliberation mechanisms (such as argumentation) is still in its infancy.

**Towards Adaptive Agents: Develop learning ability to support the eCitizen** At the architecture level, future avenues for research on learning include developing distributed models of profile management, as well as more general distributed agent learning techniques rather than just single agent learning in multi-agent domains. Other research communities have considerable expertise in the elicitation of user preferences and utilities (*marketing theory*: elicitation of desires and preferences from potential consumers; *knowledge acquisition*: learn user preferences, personalise, and construct user models; and *uncertainty in AI*: obtain expert probabilities and utilities for probabilistic belief networks), and future research in agent technologies should draw on this work. Aside from personalisation, the development of advanced technologies for personal information management raises important social issues, including privacy. Achieving truly pervasive technology to support personalisation should move society closer to the goal of universal information access, by making information accessible on the widest range of platforms in a form that is tailored to the needs of the individual.

## 2 Need and Relevance

The proposed activities will contribute to the “**e-inclusion**”, “**health**”, “**e-business and e-government**”, “**electronic and mobile commerce**”, “**eWork systems**”, “**eLearning**” and “**complex problem solving**” objectives of the sub thematic priority area of 1.1.2.i, through the provision of basic computation models and techniques. The Ambient Intelligence (AmI) vision describes an environment of potentially thousands of embedded and mobile devices (or software artefacts) interacting to support user-centred goals and activity. This suggests a component-oriented view of the world in which the artefacts are independent and distributed. Most consider autonomy, distribution, adaptation, responsiveness, etc to be key characterising features of these AmI artefacts, and in this sense they are effectively *agents*. In particular, they are likely to be function-specific (though possibly configurable to tasks) and will, of necessity, need to interact with numerous other agents in the environment around them in order to achieve their goals. Interactions will take place between pairs (in one-to-one cooperation or competition), between groups (in reaching consensus decisions), and between agents and infrastructure resources (such as large-scale information repositories, or other supporting

resources, possibly through agent encapsulation). Interactions like these enable the establishment of *electronic institutions* or *virtual organisations*, in which groups of agents come together to form larger coherent groups able to achieve some overarching goals.

The proposed activities will contribute to the “**new technologies for software and systems**” and “**control of complex distributed systems**” objectives of the sub thematic priority area of 1.1.2.ii. The need to develop infrastructure to support the kinds of systems envisaged above is well recognised, through agent-based middleware efforts, such as JADE and JAFMAS, and an emerging interest in infrastructure for Grid-enabled software. However, these are typically point solutions and new basic technologies are emerging through Web Services and related initiatives. Effective, scalable virtual organisations can only be developed and sustained through an underpinning agent infrastructure. The development of such systems requires mature development methodologies suited to the agent abstractions that facilitate take-up and deployment in the commercial world.

This infrastructure enables AmI scenarios to be realised. Artefacts or agents offering particular services can be distinguished from issues concerning *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. They can also be distinguished from issues relating to 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 artefacts and services, for example. It is particularly to ensure that large numbers of agents and services are accommodated (scalability), and that heterogeneity of agents and services is facilitated by the provision of appropriate ontologies to enable the effective interactions mentioned above.

The proposed activities will contribute to the “**knowledge technologies and digital content**” and “**intelligent interfaces**” objectives of the sub thematic priority area of 1.1.2.iv. Importantly, interactions will occur between agents in the environment, but also between agents and users, requiring greater sophistication in interface issues, and in user understanding (and modelling).

Europe has a number of research and development strengths in agent systems, including:

- Formal and logical approaches to agent systems.
- Argumentation and application to the design of agent interaction protocols
- Application to the management of complex, distributed networks, as found in the telecommunications and electricity industries. European industries have leading experience in applications of agent systems.

However, in some areas, Europe is weak relative to other regions in the world:

- Work on auction mechanisms and automated trading agents has been primarily US-led, although European researchers and companies have made major contributions.
- Work on military applications (battlefield simulations for strategy analysis; weapons and battlefield simulators for training; and agent systems for weapon control) is likely to be more advanced in the US, though information on these is not readily available.
- Game and robotics applications are led by Japan rather than the US.

At present, Europe is globally competitive in the areas of agent-based computing with some success stories and some weakness. Given the evident centrality of the paradigm and the technologies for, in particular, AmI, it seems clear that the research strengths must be supported through a period of dramatic and rapid change in the general technological environment and infrastructure. Similarly, the importance of the broad field of agent-based computing across the breadth of the sub thematic priority areas identified above suggests that positions of relative weakness, and where European activity may be enhanced, should be developed and expanded as best possible so as not to fall behind competitors elsewhere.

### 3 Excellence

Since 1998, European efforts in R&D have been supported and promoted by the AgentLink Networks of Excellence. These have supported research conferences and workshops, fostered interactions around areas of special interest, promoted agent technology to a wide industrial and commercial audience, undertaken a programme of education and training, and disseminated reports on developments in agent technology to a global audience. Through the network, the profile of European work has improved dramatically, and successful initiatives (such as the summer school) copied by other regional organisations. The networks have been instrumental in maintaining the excellence of agent-based computing in Europe by, for example, training the next generation of researchers (as others are now seeking to do) and sustaining its competitive edge, as evidenced by European success in the 2001 Trading Agents Competition (won by a German company) and disproportionate European involvement in FIPA activities, for example.

At present, the AgentLink II network of excellence for agent-based computing counts 158 full member nodes across Europe, comprising universities, research institutes and industrial and commercial organisations (62 non-university nodes). Membership criteria for assessing excellence in agent-based computing have kept the number of members down, but the network continues to grow, especially in the commercial sector. Research in Europe spans numerous areas across the range. In proposing a network of excellence for agent-based computing, we have identified a Core Group of partners involved in the coordination of the network and in the integration of activities in particular research sub-fields. Given the extent of the current network, below we limit consideration to Core partners, with their research

expertise, identified as responsible for major involvement in the main coordination and integration effort. Note that all industrial organisations have developed *research* expertise to complement related development efforts.

**Core Areas: University of Southampton, UK (Network Coordinator):** negotiation and argumentation, agent-oriented software-engineering (AOSE), Grid computing and pervasive computing. **CTU, Czech Republic:** coalition formation, security, applications in production planning, virtual enterprises, trading and humanitarian relief operations. **Motorola, France:** intelligent service aggregation and personalisation in eBusiness and mBusiness, security, trust, privacy, knowledge models. **DFKI, Germany:** application to transportation, finance, enterprise resource planning, eCommerce, control of synthetic characters in virtual worlds. **Siemens, Germany:** agent-based manufacturing. **T.U. Munich, Germany:** multi-agent learning, AOSE, socionics. **Bar-Ilan University, Israel:** automated negotiation, coalition formation. **University of Modena, Italy:** semantic integration of heterogeneous information, intelligent information agents, Semantic Web, AOSE. **University of Siena and ISTC National Research Council, Italy:** social order and social interaction (norms, conventions, roles, cooperation, power, dependence, etc.), cognitive and motivational agent architectures, agent-based social simulation. **Universiteit Utrecht, The Netherlands:** agent theories/logics, agent programming languages, agent communication and coordination languages, agent-mediated eCommerce (AMEC). **IIIA-CSIC, Spain:** AMEC, auction and negotiation mechanisms, learning and uncertainty. **Telefonica, Spain:** application to telecommunications, methodologies for agent oriented development, engineering aspects of MAS. **T.U. Catalonia (UPC), Spain:** learning in knowledge-based and agent-based systems, eCommerce, eInstitutions, application to healthcare and environmental issues. **Whitestein Technologies, Switzerland:** integration of agents with existing IT environments and platforms, mobile wireless computing, security and trust for open distributed systems. **BTexact, UK:** development, deployment and management of agent-based systems, multi-domain strategic resource management, intelligent personal assistants, agent enhanced workflow. **HP Labs Bristol, UK:** AMEC, negotiation, B2B contract management, biologically inspired computing and adaptive systems, personalisation technologies. **University of Liverpool, UK:** AOSE, argumentation, agent communications & interaction, ontologies, Semantic Web, logical models of agent systems. **Lost Wax, UK:** agent mediated eCommerce. **Qinetiq, UK:** Grid computing, Semantic Web, agent-human interaction, agent-based planning systems, application to large-scale battlefield simulations, unmanned autonomous vehicles.

As AgentLink II does at present, the network will maintain close links with related community initiatives, and undertake a coordination and integration role (as currently with Agentcities, FIPA, OMG, for example). It will be responsible for issuing calls for proposals to the community, and for distributing funding and support, and will be governed by management and advisory committees with internal and external members.

Others who have expressed an interest in participating in the wider network (possibly as Associates) include the remainder of current AgentLink II members (see <http://www.agentlink.org/members/members.html> for a full list of all 158 organisations), and particularly CETIC, Belgium; IKV++, T.U. Ilmenau, Germany; Jozef Stefan Institute, Slovenia; Fatronik, Spain; University of Cardiff, Manchester Metropolitan University, University of Salford, University of Westminster, UK.

Based on figures of 20 Core partners at 5 researchers per node, 100 Associates at 1 researcher per node, and 50 partners participating through calls for proposals, etc, we estimate an annual budget of Euro 4 million per year, over a period of 4-5 years, totalling Euro 16-20 million.

#### 4 Integration and Structuring Effect

At this critical point in the development of the technology, its take-up, and most importantly its pivotal role in contributing to broader visions of e.g., Ambient Intelligence, the Semantic Web and Grid computing, continued efforts in these respects is warranted. At the same time, an increase in relative maturity is seeing particular groups across Europe undertaking research that demands critical mass and collaboration with related groups working both in the same areas and in complementary fields. European strength in the area of agent-based computing runs broad and deep, but needs to avoid fragmentation into different camps, needs to integrate its research activities, and needs to be effectively coordinated, integrated and grown in a coherent fashion. In particular, there needs to be a focus on ensuring that the potential of agent technology, which is recognised by numerous commercial organisations, is realised to provide tangible benefits to the IT industry (through incorporation in products, development of tools and methodologies, establishment of start-ups) to the wider European R&D effort (through support for broader technological visions such as AmI and Grid computing, eliciting requirements from industry), and to establishing technological benefits for society (through agent-supported healthcare, agent mediated eCommerce, etc).

The network will seek to build on established strengths by targeting resources to ensure critical mass in particular areas and avoid over-reliance on individuals who may leave European research centres or the research sector itself. It will promote areas of particular importance to related and broader European efforts, such as AmI. It will seek to ensure coordinated research actions across multiple organisations working on the same theme, through programmes of collaborative visits, dissemination efforts to small research communities, and promotion of research results. The network will seek also to develop areas of relative weakness, by supporting efforts to grow such areas of research through similar schemes, and through small research grants that will enable researchers to gain initial expertise and will pump-prime larger research projects.

**Promotion of Results:** AgentLink has developed considerable expertise in establishing novel and accessible dissemination routes, and both formal and informal collaborative relationships with relevant bodies. The proposed network will seek to build on these existing approaches, which include: development of agent technology roadmaps for strategic direction of the field; a high quality web portal for all aspects of agent technology; high impact magazines for global distribution to the research, development and user communities, to include reports on standards activities, commercial developments, research results, etc; annual agent summer schools for academia and industry; formal interactions with related projects including Agentcities and EUTIST-AMI, and standards organizations including FIPA and OMG; monthly email updates on activities; commissioning and development of software reports; promotion of the network and support for the community through sponsorship of academic conferences; initiation of industrially focused conferences for industry. The network will also seek to promote underdeveloped regions, such as Associated States, with enhanced support.

Distinct from the more technical challenges that can be seen as clear advances in technological research or development as enumerated above, there are also important community activities that can either contribute to the success of the field's development or can prove a constraining influence. Agent-based computing is very broadly based; the connections with, and influences from, related technology and application areas, initiatives and developments with their own impetus, are critical in order to ensure better understanding, common models, better tools, etc. No single institution or country possesses anything approaching the necessary individual or corporate expertise. A sustained programme of research, development, support, promotion, coordination and integration must engage in these strategic community goals.

**Leverage underpinning work on similar problems in Computer Science:** The concepts underlying the field of agent-based computing are not unique to this particular branch of Computer Science and, as existing software technology becomes more sophisticated and moves up the application stack, it will increasingly intersect with many sub-areas of agents. For example, distributed object technologies share many similar but less sophisticated abstractions, and employ similar notions of brokers facilitating interaction between components. Similarly, existing approaches to software engineering have much to offer the current efforts to develop agent-specific development methodologies. More generally, theories of agent interaction have yet to draw in detail on abstract theories of distributed computation. It is important that stronger links with these traditional areas of computing are established and reinforced in order to leverage their underpinning work, while also providing scope for work in the development of agent-based systems correspondingly to inform them.

**Link with related areas in Computer Science dealing with different problems:** Related disciplines that typically address different problems to those tackled by agent-based computing can also provide valuable inputs to R&D. These include the fields of artificial life, computational biology and computational economics, which have developed an arsenal of techniques that may be appropriate for application to agent-related problems. The building of bridges to different disciplines should thus also focus on areas with different target problems, so that techniques may be adopted wholesale, or in hybrid approaches to offer new solutions. This can be clearly observed in the potential of artificial life techniques for agent-based simulation, and learning and adaptation in such contexts, as well as work on Uncertainty in AI.

**Extend and deepen links with other disciplines:** For many disciplines, the links to agents are already deep and fruitful, such as those with economics, logic and philosophy. For other disciplines, there exists potential to develop closer connections. For instance, agent systems designers could draw more extensively on political theory and sociology in the design of agent societies, and on decision theory in the assessment of performance of agent systems. Similarly, agent-mediated eCommerce has yet to make great use of the models and techniques developed by marketing theorists for preference elicitation or the diffusion of innovations.

**Encourage industry take-up:** Commercial deployment of agent systems is currently confined to early adopters in some segments of industry and government (e.g., utility companies) and agent systems have yet to achieve widespread deployment, possibly through immaturity of the technology. This can be addressed through working commercial prototypes for specific industry sectors, made available for commercial use. There may also be a lack of awareness of the potential applications of agent systems, requiring early adopter case studies to be prepared, both successful and unsuccessful, with an analysis of the reasons for success or failure. Such case studies should also include assessments of the resources and timescales required, and the factors critical to their successful deployment. Additionally, high deployment costs are a feature of any new technology, and this is likely to be true for agent systems. As agent design tools and standard methodologies are developed, and as development teams gain greater experience, these costs should fall. To ensure that these experiences are disseminated beyond the early adopter community to other organizations, it is vital that best practices for agent-oriented development and deployment be identified and publicised.