Europe's IST-funded Network of Excellence for Agent-based Computing

ISSN 1465-3842

Issue 11

December 2002

AgentLink news

Features

Agent Technology: Roadmap Abstract and FP6 Expression of Interest Towards AgentLink III?

FIPA and the evolution of the specifications

The **Agentcities** Task Force

Agents and Logic: an Inter Network SIG

Book Reviews

An Introduction to Multiagent Systems

Understanding Agent Systems

Special Interest Groups

Agents that Learn, Adapt and Discover (ALAD)

Conference Reports

NETTAB 2002 workshop on Agents in Bioinformatics

Fifth International Workshop on Cooperative Information Agents (CIA 2001)

The 2nd International Workshop of Central and Eastern Europe on Multi-Agent Systems (CEEMAS)

The Second Symposium on Adaptive Agents and Multi-Agent Systems (AAMAS-II)

UK Special Interest Group on Multi-Agent Systems (UKMAS)

The Fourth European Agent Systems Summer School (EASSS 2002)

Contents

With this issue of AgentLink News, we mark the handover of the editorial reins from Paul Davidsson to Andrew Byde. After four years in the job, Paul has decided to concentrate on his role coordinating the Agent-Based Social Simulation SIG, and although we are sad to lose him, he will remain on the newsletter team. Paul has done an excellent job in this time, but we have the best possible replacement in Andrew Byde of HP Labs, who is taking over as editor, starting with this issue.

As we move into the final year of AgentLink II, attention is turning towards the future development of the field of agent-based computing and its role in the Sixth Framework. In light of this, the first article of this issue is a revised version of the Expression of Interest submitted on behalf of AgentLink to the European Commission earlier in 2002, and represents the start of efforts to coordinate the community's activity.

AgentLink is keen to support the community in FP6, both by preparing a proposal for a Network of Excellence, and by providing an information resource on related efforts in the field. The web pages at **http://www.agentlink.org/fp6/** provide one way of doing that. If you have other ideas or suggestions, mail the AgentLink Coordinator, Michael Luck as mml@ecs.soton.ac.uk

Michael Luck
University of Southampton

FEATURES

Book Reviews

Conference Reports

Special Interest Groups

Agent Technology: Roadmap Abstract and FP6 Expression of Interest Towards AgentLink III?

Michael Luck
University of Southampton
United Kingdom
mml@ecs.soton.ac.uk

As the European Commission's Sixth Framework Programme approaches, many in the community are considering the submission of proposals in the area of agent-based computing. AgentLink has been involved in two particularly important developments in relation to this: it submitted an Expression of Interest in support of a new Network of Excellence; and it has developed a draft Agent Technology Roadmap that seeks to set the agenda for the next ten years of agent research. This article summarises both.

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 require further development of 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 had some success, e.g. with Agent UML, 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, 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. As more sophisticated interactions become common, open systems development will also require libraries of interaction protocols designed for specific interactions. These may use existing agent communication languages, as various auction protocols do 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 be. 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 begun 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, namely

- *marketing theory*, concerning the elicitation of desires and preferences from potential consumers;
- *knowledge acquisition*, concerning the learning of user preferences, personalisation, and construction of user models; and

- *uncertainty in AI*, concerning the acquisition and use of expert probabilities and utilities for probabilistic belief networks.

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 "**e-inclusion**", "**health**", "**e-business and e-government**", "**electronic and mobile commerce**", "**eWork systems**", "**eLearning**" and "**complex problem solving**" 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 and 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 also contribute to "**new technologies for software and systems**" and "**control of complex distributed systems**". 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, for example, the physical infrastructure needed to support effective interaction through sensors and actuators, and the physical connectivity for supporting quick and efficient interactions. They can be 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. It is particularly important 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.

Finally, the proposed activities will contribute to "**knowledge technologies and digital content**" and "**intelligent interfaces**". 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; and
- 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 areas of relative weakness, and areas where European activity may be enhanced, should be developed and expanded as much as 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) have been 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 demonstrated by, for example, European success in the 2001 Trading Agents Comptition (won by a German company) and disproportionate European involvement in FIPA activities.

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. A new network of excellence for agent-based computing will require the integration of activities in particular research sub-fields as follows.

> *"The profile of European agent work has improved dramatically"*

Core Fields: negotiation, argumentation, agent-oriented software-engineering (AOSE), Grid computing, pervasive computing, coalition formation, security, applications in production planning, virtual enterprises, trading and humanitarian relief operations, intelligent service aggregation and personalisation in eBusiness and mBusiness, security, trust, privacy, knowledge models, application to transportation, finance, enterprise resource planning,, control of synthetic characters in virtual worlds, agent-based manufacturing, multi-agent learning, socionics, coalition formation, semantic integration of heterogeneous information, intelligent information agents, Semantic Web, social order and social interaction (norms, conventions, roles, cooperation, power, dependence, etc.), cognitive and motivational agent architectures, agent-based social simulation, agent theories/logics, agent programming languages, agent communication and coordination languages, agent-mediated eCommerce (AMEC), auction mechanisms, uncertainty, application to telecommunications, methodologies for agent oriented development, engineering aspects of MAS, eInstitutions, application to healthcare and environmental issues, integration of agents with existing IT environments and platforms, mobile wireless computing, trust for open distributed systems, development, deployment and management of agent-based systems, multi-domain strategic resource management, intelligent personal assistants, agent enhanced workflow, B2B contract management, biologically inspired computing and adaptive systems, personalisation technologies, agent interaction, ontologies, agent-human interaction, agent-based planning systems, application to large-scale battlefield simulations, unmanned autonomous vehicles.

Any new network will need to maintain close links with related community initiatives (as AgentLink II does at present), and undertake a coordination and integration role (as AgentLink II does with Agentcities and FIPA, 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.

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 such as 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 undertake research that demands critical mass and collaboration with related groups, both those working in the same areas, and those working 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);
- the wider European R&D effort (through support for broader technological visions such as AmI and Grid computing, eliciting requirements from industry); and
- society at large (through agent-supported healthcare, agent mediated eCommerce, etc.).

A new 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 also seek 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. At the same time, such links can also allow work in the development of agent-based systems to inform these traditional areas of computer science.

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.

5 Acknowledgements

This is an edited version of the Expression of Interest submitted by the AgentLink II coordinator on behalf of the network to the European Commission's IST programme. It draws heavily on an early version of AgentLink's Agent Technology Roadmap, written and collated by Michael Luck, Peter McBurney and Chris Preist, a public draft of which is available at www.agentlink.org/roadmap. Comments on both this EoI, and particularly the Roadmap, are invited from AgentLink members and from the wider community.

Although the roadmap is unlikely to satisfy everyone, it seeks to provide a vision and guide for as broad a constituency as possible. It is a living document and relies on constructive criticism and input. Vague suggestions and comments are difficult to address; please mail detailed comments to Michael Luck at mml@ecs.soton.ac.uk

Michael Luck of the University of Southampton will coordinate a full proposal for a new Network of Excellence in the area of agent-based computing. The community will be kept abreast of developments through the AgentLink Email Update list. If you do not receive the regular email updates, but would like to do so, mail coordinator@agentlink.org.

Draft

Agent Technology Roadmap

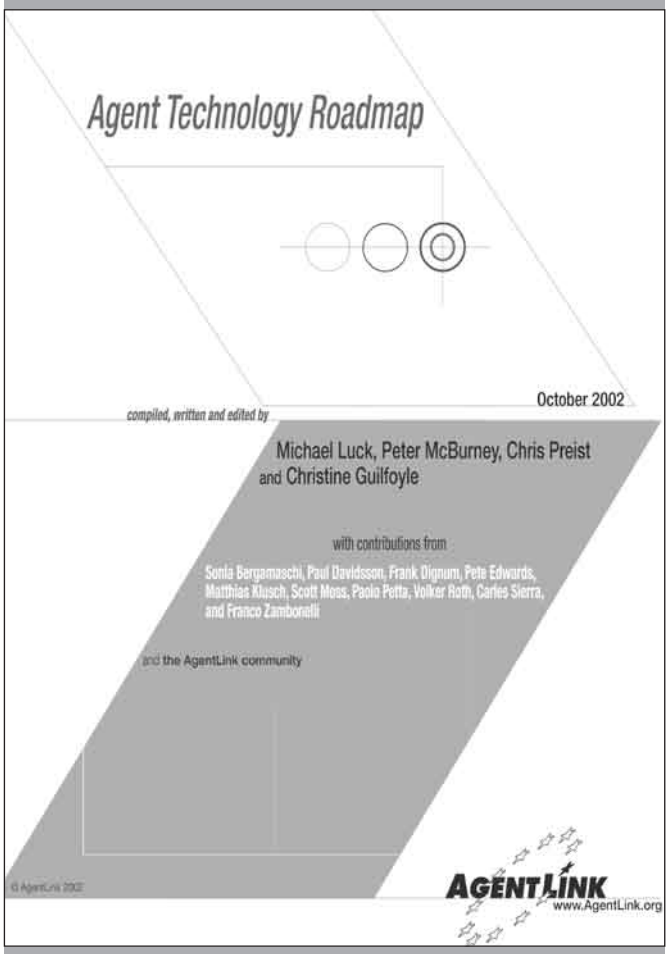

available from:

www.agentlink.org/roadmap/

february 5th Spring School 2003

easss

The 5th European Agent Systems Spring School

After successful Summer Schools on Agent Systems in Utrecht, Saarbruecken, Prague and Bologna, Agentlink is organising the fifth such event in Barcelona, Spain, 10-14 February, 2003. This is the week after AgentLink SIG Meetings and Agentcities Information Days take place there.

EASSS 2003 will consist of a mixture of introductory and advanced courses delivered by internationally leading experts in the agent field, and will cover the full range of theoretical and practical aspects of agent-based computing.

It is open to anyone from research or industry, both AgentLink members and non-members alike. A registration fee will be charged to cover costs, but some support will also be available for PhD students.

Details available from:

www.agentlink.org/happenings/easss/2003/

FIPA and the evolution of the specifications

Monique Calisti
Whitestein Technologies AG
Switzerland
mca@whitestein.com

The core mission of the FIPA (Foundation for Intelligent Physical Agents) standards consortium is to facilitate the interoperation of agents and agent systems across multiple vendors' platforms. In the increasingly open and distributed agent-based environment we see emerging nowadays, the need for standard mechanisms and specifications is crucial for ensuring interoperability of distinct autonomous systems. For this purpose, FIPA has been working since 1997 on specifications that range from architectures to support agents' communicating with each other, communication languages and content languages for expressing those messages, and interaction protocols to expand the scope from single messages to complete transactions. The core message of FIPA is that through a combination of speech acts, predicate logic and public ontologies, we can offer standard methods of interpreting communication between agents in a way that respects the intended meaning of the communication. The FIPA community has been very active from several different perspectives:

- Developing, providing and maintaining a set of standard specifications.
- Defining guidelines for industrial development of agent platforms and agent-based applications (currently, there exist several FIPA compliant platforms that are widely used in several contexts. More information is available on the FIPA web site, http://www.fipa.org).
- Providing an important discussion forum:
  - FIPA meetings (open and free) enable industrial players and researchers from both the industrial and the academic world to come and work together. The next FIPA meeting will be held the 10-11-12 February 2003 in Palermo, Italy.
  - On line discussions through specific and/or more generic groups (e.g., chat@fipa.org, see also the FIPA web site for more information about the current available mailing lists).
- Ensuring strong liaison with universities, projects (including AgentLink and Agentcities), standards (such as W3C, OASIS, IETF, OMG, AUML, etc), and implementations.
- Sponsoring relevant events in the agent world and distributing the main results in several ways (FIPA Inform! newsletter, web site, white papers, etc.)

When an idea for potential FIPA work has been formed, elaborated, and discussed (chat@fipa.org or specific existing working groups), it is submitted as a work plan (clear timeline and committed participants) to the FIPA Architecture Board (fab@fipa.org). A specific work plan can be carried out by a Technical Committee (TC) for normative specifications or by a Working Group (WG) for informative specs, applications, field trials, etc. Specifications are then created:

- Preliminary (P): Draft under discussion (TC approval),
- Experimental (X): Stable, suitable for implementation (FAB approval),
- Standard (S): Stable, successfully implemented (FAB and membership approval),
- Deprecated (D): Potentially unnecessary (FAB and membership approval),
- Obsolete (O): Rendered unnecessary.

Over the last few years, FIPA has been playing a unique role in bringing together various organizations (both industrial and academic) and key players from several domains such as the IT world, the Telecommunication domain, etc. However, given the fast technological changes in modern communication networks and distributed computational systems, the strong influence of the Internet explosion and the pressing needs for more flexible and dynamic ways of providing services to an increasing and varied number of potential customers, FIPA's focus and its specifications have evolved and adapted to these changes. Currently, FIPA's main activity is focusing on:

- Promoting to 'standard' status a core set of FIPA specifications by the end of 2002. The X2S Technical Committee has been very active during the whole year and within the next few weeks the FIPA community will be invited to vote for 35 specs (www.fipa.org/activities/experimental_to_standard.html).
- Working on a new semantic framework to reflect the needs of verifiability and conformance. In particular, the objective is to adopt or define a semantic framework that can give an account of FIPA's existing communicative acts and interaction protocols as well as a number of additional constructs.
- Creating new specifications to ensure that interoperability between FIPA-compliant agent platforms and/or platform fragments can be maintained in ad hoc networks.
- Standardising methods for knowledge sharing and filtering through ontological representations that allow interoperating systems to automate message processing with respect to cross-referenced semantic classification.
- Building a service model for representing, modelling, discovering and using services. This work is intended to provide a way to interface agent-based services and traditional web services.
- Finally, creating new specifications for providing security mechanisms and support in FIPA compliant agent systems.

Besides these main flows of action, a number of people are also considering planning and scheduling issues to promote and support the development of agent applications that require coordination of task execution, as well as work on AUML and agent methodologies.

The most effective way for both members and non-members of FIPA to contribute to the standardization process, is to make use of the existing specifications and report back to the community, to participate in the FIPA meetings (open and free!) and/or the mail discussions, and to make an active contribution to the work and to the creation of the various documents. This is why FIPA strongly believes in the value of close interaction with an active, dynamic and wide community such as AgentLink, and encourages collaborations and synergies at various levels as well as projects in the context of the 6th EU research framework.

Acknowledgments: many thanks to the whole FIPA community and in particular to the FIPA board of directors for their contribution to this overview about FIPA.

The Agentcities Task Force

"...creating an open world of intelligent automated services..."

Jonathan Dale
Fujitsu
USA
jdale@fla.fujitsu.com

Steven Willmott
Universitat Politecnica de Catalunya
Spain
steve@lsi.upc.es

Bernard Burg
Hewlett-Packard
USA
bernard_burg@hp.com

Patricia Charlton
Motorola
France
charlton@motorola.com

Paul O'Brien
British Telecommunications
UK
paul.d.obrien@bt.com

Mission

One of the greatest technology challenges facing both research and business today is realising the vision of open networks populated by automated service components interoperating seamlessly to achieve user and business goals. Achieving this vision requires the integration of many technologies and incremental evolution of deployed systems through experimentation and experience with deployed systems.

The objective of the Agentcities Task Force is to facilitate creation, exploitation and use of a worldwide, public, open, heterogeneous and interoperable environment in which autonomous services can be defined, deployed and utilised in a dynamic, composable and value added way. The network environment will act as a testing ground for cutting edge technologies and applications as well as a shared environment to support collaborative work between leading research and industry organisations.

The infrastructure will build on the existing Agentcities Network, which already involves more than 100 organisations worldwide, with 50 experimental service platforms deployed. The network structure will draw upon existing standards from Web Services, GRID computing and e-Business worlds combined with cutting edge research technologies such as Intelligent Agents and Semantic Web technologies.

The goals of the ACTF include:

- Supporting technology development, take-up, awareness and standardisation through activities such as technical meetings, workshops, conferences and other types of promotion,
- Promoting and supporting the deployment of shared commercial and non-commercial infrastructure, software and services to support the development of the network environment, and,
- Interacting extensively with standards bodies in the field including FIPA, W3C, IETF, GRID forum, OASIS and the WS-I as well as developing de-facto and de-jure technical recommendations (standards) as necessary to promote interoperability, usage and evolution of the network environment

The role of the Agentcities Task Force will be to act as a focal point for activities in the Agentcities Network, supporting collaboration between related projects worldwide, helping to establish consensus on technical issues, and delivering results to appropriate standards bodies.

Formation

ACTF will be incorporated as a non-profit organisation under UK law in the form of a "Limited Company Limited by Guarantee" operated by its members through equal voting representation at general and other assemblies. Membership is open to any organisation.The responsibilities of the ACTF include:

- Carrying work out under an IPR framework based on the IETF's model ensuring free, fair and unrestricted access to all technical results to both members and third parties,
- Ensuring that all funds will be used to further the objectives of the organisation on behalf of its members, and,
- Working closely with leading standards organisations such as IETF, W3C, GRID Forum, the WS-I and FIPA to ensure that technical experiences gained in the Agentcities network benefit the development of prevailing standards.

The draft statutes are currently being ratified by legal experts to ensure correctness under UK law and we expect the formation of the organisation itself to take place March, 2003.

Following this we expect to be able to issue a call for membership in April, 2003; after this point members will be able to join at any time.

Costs and Benefits

ACTF is a not-for-profit organisation and all its funds will be used to further the common objectives of its members through the organisation of conferences, meetings, hosting of web and other services, promotion and liaison with relevant standards bodies in the field. The primary mechanism for raising funds will be through annual membership dues:

University or non-profit organisation	500 €/year
Small or medium enterprise (<100 staff)	1500 €/year
Large corporation (>101 staff)	3000 €/year

Members will benefit from high visibility as a supporter of the initiative, a strong voice in setting the strategy and direction of the organisation (through voting rights at general assemblies, eligibility for board of directors, advisory board, working group chair and other leading positions in the organisation). As a global consortium of the leading Universities and Companies in the field, ACTF will enable collaborative links between members and create sufficient critical mass to influence technology development in the key standards bodies.

Organisations joining in the first wave of members will be additionally recognised as founding members of ACTF.

About Agentcities

Agentcities is an initiative that was first conceived in January 2000 to create a next generation Internet that is based upon a worldwide network of services that use the metaphor of a real or a virtual city to cluster services. These services, ranging from eCommerce to integrating business processes into a virtual organization, can be accessed across the Internet, and have an explicit representation of the capabilities that they offer. The ultimate aim is to enable the dynamic, intelligent and autonomous composition of services to achieve user and business goals, thereby creating compound services to address changing needs. This is in direct contrast with the current Internet, where users have to interact directly and constantly with software, that is, Web browsers, to achieve their goals and fulfil their requirements. The initiative will build on a wealth of innovative technologies including agent technology, Semantic Web technologies, UDDI discovery services, eBusiness standards and Grid Computing. Application areas already envisaged range from eHealth and eLearning to manufacturing control, digital libraries, travel and entertainment services.

More information is available from :
http://www.agentcities.org/ or email actf@agentcities.org

AGENTS AND LOGIC:
AN INTER NETWORK SIG

Michael Fisher
University of Liverpool
United Kingdom
M.Fisher@csc.liv.ac.uk

This article provides a brief overview of a new activity, funded jointly by AgentLink and CoLogNET, that aims to enhance interactions between the Agent and Computational Logic communities. We begin by providing an overview of the two networks involved (although readers should refer to the relevant websites for full information).

AgentLink

AgentLink [http://www.agentlink.org] is the EU's Network of Excellence for agent-based computing. AgentLink coordinates research and development activities in agent-based systems and supports a range of activities aimed at raising the profile, quality, and industrial relevance of agent systems research and development in Europe.

The view of agents from AgentLink is that "an agent is an autonomous software system: a system that can decide for itself what it needs to do."

In addition, agents are deemed important for several reasons:

- they are seen as a natural metaphor for conceptualising and building a wide range of complex computer systems;
- they cut across a wide range of different technology and application areas, including telecoms, human-computer interfaces, and distributed systems;
- they are seen as a natural development in the search for ever-more powerful abstractions with which to build systems.

AgentLink divides its activities into five main areas:

1. Industrial action ("facilitating technology transfer through a program of industrial meetings, workshops, standardisation updates, and working groups");
2. Research coordination ("promoting excellence in European agent research through support for workshops, special interest groups, and dissemination of research results");
3. Teaching and training ("establishing agent related skills throughout Europe by support for summer schools and courses");
4. Special Interest Groups ("focusing on the development of communities around areas of strategic importance and providing input to the management committee from the SIG members, as well as developing the technological roadmap");
5. Information Infrastructure ("providing an infrastructure through which AgentLink can do its work, including a website, regular newsletter, email list, and an awareness programme").

AgentLink is very successful and very large. If you are interested in becoming a member of AgentLink, see http://www.agentlink.org/members/join-al2.html

CoLogNET

CoLogNET [http://www.colognet.org] is a new EU Network of Excellence that builds upon the work of Compulog Net, the first European Network of Excellence for Computational Logic, that provided the role model for many networks to follow. However, CoLogNET has a much broader scope than Compulog Net, taking into account the recent growing diversity of Computational Logic and logical applications.
CoLogNET aims to "help unify and integrate the separate sub communities and is therefore built around an entirely new and broader structure representing the new communities involved as well as some of the older ones."

An ambitious, longer term, goal of the network is to promote Computational Logic as an academic discipline in its own right, on a par with Mathematics or Physics. To this end, the network will provide a European supporting infrastructure and information sources for this new academic discipline and its many potential application areas.

CoLogNET activities are planned in nine main areas:

1. Logic Methodology and Foundational Tools;
2. Logic and Constraint Programming;
3. Computational Logic Systems and Environments, Component Based Software Development;
4. Logic-based Data and Knowledge Systems;
5. Logic and Multi-Agent Systems;
6. Logic and Natural Language Processing;
7. Formal Methods, Specification and Verification;
8. Automated Reasoning, Deduction, Theorem Proving and Model-Checking;
9. Logic in Mechanical and Electrical Engineering.

AgentLink/CoLogNET joint activity on "Agents and Logic"

Formally, this activity is one of the main areas within CoLogNET (area 5, as above) and is a special interest group within AgentLink. Both networks, showing their commitment to this important area, fund the activity.

Background:

The link between agents and logic is important to both the Agents and Computational Logic communities: to the Agents community since logic-based approaches already have a strong profile, and are likely to grow in importance especially as verification becomes more necessary; to the Computational Logic community since agents represent a key application area involving a wide range of relevant techniques. While CoLogNET on its own might contain a little about agent applications, and AgentLink contain

a little about logic, this formalised activity aims to allow more expertise/ideas to flow between the networks, thus providing more chance of productive technology transfer.

Thus, in the future, this activity will contribute to the uptake of logic-based agent technology and logical methods for agent-based systems. This, in turn, may impact on forthcoming standards, for example where semantics/verification/compliance is an issue.

Objectives:

- to disseminate the state of the art in key areas concerning logic and agents throughout both communities, and industry, via dedicated symposia;
- where possible to disseminate key breakthroughs in this area to wider academic/industrial communities by developing publications examining these specific areas;
- to enhance collaboration between groups, networks, academia/industry, etc;
- to showcase successful applications concerning the verification and implementation of agents using logic-based approaches;
- to identify key future research issues and key requirements for wider use of these techniques.

Activities:

The work of this area/interest group is centred around two symposia, one on "logic-based agent verification", the other on "logic-based agent implementation", both of which provide very strong links between the networks. Each symposium will aim to consolidate expertise and stimulate research/collaboration in each area, and to provide input into both standards efforts and the development of technological roadmaps for the area. They will also endeavour to showcase successful applications and enhance the uptake of logic-based agent technology. Via industrial presentations/sessions, they will also attempt to elicit key problems in industry that could be tackled with a logic-based approach, and to identify barriers to such technology transfer.

It is planned that each symposium be co-located with a general AgentLink or CoLogNET supported event. For each symposium, funds will be provided for one invited speaker from outside the EU and for the travel/accommodation support for a number of other participants.

Note that the aim of these symposia is not to compete with existing logic-agent research workshops, but to provide a forum whereby the state of the art within a particular topic concerning logic and agents can be reported, examined and transferred not only to both CoLogNET and AgentLink communities, but to industry and the wider academic community.

As part of this transfer activity, we plan to provide a web-accessible focal point for research, education and applications of "logic and agents", building up a set of working exemplars concerning the role of logic in agent-based systems. This will also provide wider access to tools for (and tutorials/explanations of) verification, logic-based programming, etc, of agent-based systems.
Finally, the aim is that each symposium should be co-organised by a leading EU researcher in the topic area. Thus, for example, the symposium on "logic-based agent implementation" will be co-organised with Juergen Dix [http://www.cs.man.ac.uk/~jdix]

Plan:

In broad outline, the planned activity is as follows

[late 2002] Initial WWW repository
[2003] Symposium on "logic-based agent implementation"
[2004] Symposium on "logic-based agent verification"

An Introduction to Multiagent Systems

A Review by **Gerhard Weiss**

Technische Universität München
Germany
weissg@in.tum.de

Michael Wooldridge
John Wiley & Sons Ltd, 2002
ISBN 0-471-49691-X

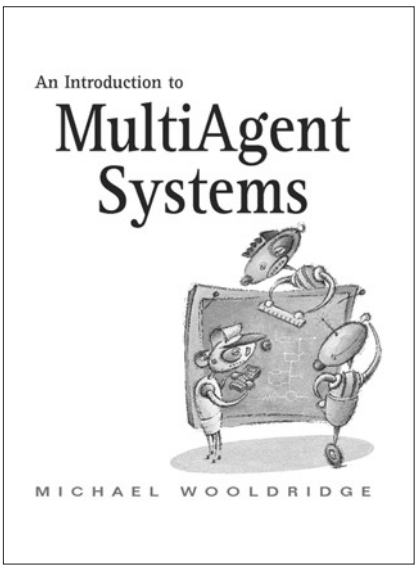

Multi-agent systems are systems composed of computational entities, known as agents, that act and interact flexibly and at least to some extent autonomously. Since its inception in the 1970s, the field of multi-agent systems has developed rapidly. In fact, compared to its relatively young age the field has already achieved a remarkable level of maturity. Today multi-agent systems are not merely an established subject of worldwide research, but are on the point of becoming a technological key ingredient for an increasing number of industrial, commercial and scientific applications in complex (open, networked, dynamic, etc.) domains such as e/m-commerce, information and systems management, peer-to-peer computing, and social simulation. Because of the field's importance, popularity, and the respectable amount of knowledge and experience accumulated over the past 30 years, there is an urgent need for a profound textbook that makes the field easily accessible to newcomers and readers having little previous knowledge of multi-agent systems. Now the first book is available which explicitly aims at fulfilling this demand, and which does so very successfully: "An introduction to multiagent systems", by Michael Wooldridge, a professor of computer science at the University of Liverpool and a leading authority in the field who has made many tremendous contributions himself in the past ten years or so. This 350-page book, which is specifically intended for undergraduates and beginning graduates of computing or information technology, has several features which make it a first-grade textbook for introductory courses on multi-agent systems. These features are, briefly, that: the book covers key topics in the field broadly, including a variety of elementary concepts, principles, and methods; it contains pointers to advanced reading and exercises of varying degrees of difficulty; and it is self-contained, coherent, and written in an illustrative and motivating style. Some students will be pleased to hear that the book avoids any unnecessary formalism, whereas necessary formalism is always accompanied by sufficient explanation and examples. A further useful feature is that the book comes with a complete set of slides for use in courses. These slides, and other material (e.g., an extensive list of agent-related web links), are available at the book's dedicated web page www.csc.liv.ac.uk/~mjw/pubs/imas/.

The book is divided into four main parts, and an interesting appendix that gives a history of the field. The first part (consisting of chapter 1) introduces the readers to the field. The second part (chapters 2 to 5) is about major types of individual agents, differing in their architectures and in the mechanisms and techniques upon which their decision making and/or reactive behaviour is based. The third part (chapters 6 to 10) is about major types of interaction in collections of agents, differing in the mechanisms and techniques that enable the involved and possibly self-interested agents to achieve consensus and behave as a coherent whole. Moreover, this part overviews the main methodologies for multi-agent system development. The fourth part (chapters 11 and 12) offers supplemental material, focusing on applications of multi-agent systems and on formal approaches to reasoning about and specifying multi-agent systems. Each chapter begins with motivational remarks, and ends with two sections which are particularly desirable for a textbook: a section providing pointers to sources of advanced reading (including a "class reading suggestion"); and a section providing exercises (including "class discussion suggestions") which are graded on a scale from one ("basic comprehension level") to four ("research project level").

A brief overview of each chapter follows. Chapter 1 ("Introduction") motivates the field by describing its vision and different ways of viewing multi-agent systems. The chapter also discusses the most commonly voiced objections to multi-agent systems. Chapter 2 ("Intelligent Agents") introduces the concept of a computational agent and discusses the main properties usually ascribed to an agent. Chapter 3 ("Deductive Reasoning Agents") deals with agents that logically reason about their environment through proving theorems. The chapter shows how logic-based formalisms can be employed to specify multi-agent systems, and

describes two famous representatives of such formalisms: AGENT0, a prototypical language for programming agents in terms of mentalistic notions such as belief and desire, and Concurrent MetateM, an agent specification language based on temporal logic. Chapter 4 ("Practical Reasoning Agents") concentrates on agents that have intentions, plan ahead, and thus reason and decide more like humans do. The chapter describes two approaches to practical reasoning agents: HOMER, a planning agent with linguistic abilities, and PRS, perhaps the most famous belief-desire-intention architecture for agents. Chapter 5 ("Reactive and Hybrid Approaches") deals with agents that react to their environment without explicitly reasoning about it. The best-known architecture for such agents, which is also described in this chapter, is the subsumption architecture. Furthermore, the chapter describes different approaches to building agents that both plan and react, and presents TouringMachines and InteRRaP as two illustrative examples of such hybrid agents. Chapter 6 ("Multiagent Interactions") introduces basic concepts related to agent-agent interaction such as utility and preference, dominant strategy and Nash equilibrium, and competitive and zero-sum interaction. Chapter 7 ("Reaching Agreements") concentrates on mechanisms that enable agents with divergent interests and conflictive goals to achieve consensus. These mechanisms are automated auctioning, negotiation, and argumentation. Chapter 8 ("Communication") treats fundamentals of agent-agent communication. The chapter familiarises the reader with the notion of speech acts and with agent communication languages (KIF, KQML, and FIPA ACL), and also points at the important role of ontologies and coordination languages. Chapter 9 ("Working Together") shows how multiple agents can act jointly and in a concerted manner. Among the key issues treated in this chapter are cooperative problem solving, task and result sharing, multi-agent planning, and various forms of coordination (e.g., by mutual modeling, by norms, and through joint intentions). Chapter 10 ("Methodologies") surveys several general analysis/design and specification methods for agent-oriented systems, including the AAII method, Gaia, Agent UML, DESIRE, Cassiopeia, and Z-based agent specification. It also overviews main aspects of and languages (Telescript and Tcl) for mobile agent construction, and describes a number of pitfalls of agent-oriented development. Chapter 11 ("Applications") lists a number of available agent applications in a variety of domains, including workflow and business process management, distributed sensing, electronic commerce, human-computer interfaces, and many others. Finally, chapter 12 ("Logics for Multiagent Systems") focuses on logic-based approaches to multi-agent systems specification. Among other things, this chapter motivates the use of modal logics for reasoning about agent systems and discusses the value of formal methods in agent-oriented software engineering.

Generally, an introductory textbook should cover topics that are both at the heart of the field and understood in sufficient depth. To identify the topics in the multi-agent systems field that meet this criterion is quite challenging for two major reasons. First, many topics in this vital field are still active research topics, and this makes it difficult to assess precisely their developmental stage. Second, the field covers a broad, multidisciplinary, and not yet fixed range of themes and issues; this makes it difficult to unambiguously determine the field's centre. Wooldridge successfully manages this challenge. His book offers a balanced choice of topics, which together give a good overall picture of the field and, in particular, reflect the field's concern with both individual agents and collections of agents. A topical exclusion that I consider noticeable is multi-agent (cooperative/competitive) learning.

The features of this book that I find most valuable are the pointers to advanced reading, the exercises, and the accompanying slides. The pointers to advanced reading are well suited for guiding interested students in further exploration of the field; teachers who use this book but want to discuss some aspect in more detail will also find these pointers useful. The exercises allow students to test their understanding of the material, and teachers may use them as a starting point for creating their own exercises. The exercises of levels one, two and, to some extent, three, should be fully solvable on the basis of the material provided in the book, whereas level-four exercises require a reader to study additional material in order to find reasonable solutions. However, it would have been great if more exercises were included. What is not available to date is a solution manual for the exercises. (Two sample exam questions are available at the book's web page.) The lecture slides contain all the book's key statements and thus are of value to every teacher who uses this textbook. Moreover, I found the slides very helpful for presenting selected background information in a one-semester course that presupposed knowledge about multi-agent systems.

In conclusion, this is an excellent introductory textbook, which fills a significant gap in the available literature on multi-agent systems. It gives novices a good feeling for, and a basic understanding of, key issues in the field, and this makes it a must for all students having little or no background in multi-agent systems and artificial intelligence.

Gerhard Weiss is the editor of *Multiagent Systems*, 1999, MIT Press.

Understanding Agent Systems

A Review by **Simon Parsons**
City University of New York
USA
parsons@sci.brooklyn.cuny.edu

Mark d'Inverno and **Michael Luck**
Springer Verlag, 2001
ISBN 3-540-41975-6

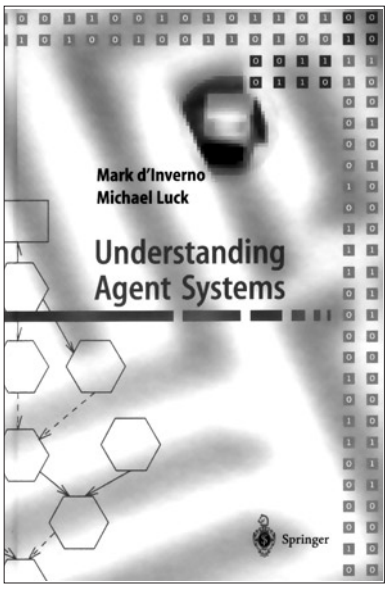

There are two main ways that natural science proceeds - theory and experiment. Some scientists create theories and use them to make predictions about how they think the world should work. Other scientists carry out experiments, measuring precisely how the world does in fact work. Theoretical and experimental works are symbiotic. Theories start as attempts to explain the results of previous experiments, and, once the old experiments are explained, are then used to make predictions that have not yet been tested. These predictions then suggest new experiments, to test the predictions, and the results of the experiments either cause theories to be revised or provoke new predictions.

Computer science is rather different. In most of computer science, the phenomena that we are studying are artificial. They are human constructs, typically algorithms in one form or another. As a result, while computer scientists certainly do theoretical and experimental work, the feedback between theory and experiment is rarely as tight as in the natural sciences. A typical experiment, instead of testing if a theory gives a correct prediction about the world, tests whether an algorithm gives a better performance on some test problem than another algorithm. Since there is no general way to rank the test problems, experimental work in computer science lacks the gold standard that the real world provides for natural sciences.

Now, this should not be read as an attack on experimental work in computer science (something that I am a firm believer in[1]). There simply is no gold standard, and we have to do the best we can with what we have got. However, it does have an impact on theoretical work. Without the drive to explain experimental results, theoretical work in computer science, and in particular in artificial intelligence, is often some distance from reality (even from a reality, made up of notions such as computation and intelligence, that is somewhat abstract to begin with). As a result, computer science theory often seems rather unconnected, existing only in relation to other theory - and rather like J. M. Barrie's fairies, apt to disappear if people stop believing in it.

To survive, computer science theory has to be relevant to what computer science practitioners (who stand in the same relation to computer science as engineers do to the natural sciences) are doing. Theory has to illuminate the work of practitioners, identify good ways of doing things, and rule out bad ways of doing things. Computer science theory has to help practitioners - all those theorems have to have a use, and preferably one that saves time and money.

One of the best ways that computer science theory can, in my view, be relevant and useful is to clarify what practitioners are doing. The practice of computer science can often be difficult and confusing, and a little clarity can help a lot (just think of all the debugging hours saved by the advent of structured programming following "GOTO considered harmful", or the move from assembly code to high level languages). While "Understanding agent systems" is unlikely to have the same kind of impact as structured programming, it makes a contribution in a similar vein.

In particular what "Understanding agent systems" presents is an approach to the formal specification of agent-based systems using Z. Starting right at the lowest level, with the idea of entities (collections of attributes), the book builds a consistent picture which integrates many of the features discussed in the literature of agents and multi-agent systems, right up to the level of co-ordination techniques (such as the contract net and social dependence networks). It is undoubtedly a clear and most comprehensive attempt to describe agent-based systems in a unified manner, and if the text isn't quite as lyrical as the first author's piano technique, well that is understandable. What the book provides is a view that practitioners can pick up and use without having to wade

Notes

1 To misquote Oscar Wilde, I could watch people doing good experimental work all day.

through the murky complexities of the rest of the literature. In that sense "Understanding agent systems" is theoretical computer science of the most helpful kind.

The only problem with this approach, which is inherent in any such formalisation, is that it picks a very specific view to formalise. This, of course, means that if one disagrees with some nuance of the formalisation, say the idea that a inanimate object such as a cup can be an agent, there is no way to use the formalisation as it stands. One would have to recreate it all again. This, however, seems much less likely to worry practitioners looking for a guide through the literature than other theorists who have strong opinions on the agenthood or otherwise of their favourite coffee cup, and does not significantly detract from the strength of the book.

The 2nd International Workshop of Central and Eastern Europe on **Multi-Agent Systems** (CEEMAS)

September 2001

Barbara Dunin-Keplicz
Polish Academy of Sciences
Warsaw University
Poland
keplicz@mimuw.edu.pl

The Second International Workshop of Central and Eastern Europe on Multi-Agent Systems was held in Cracow, Poland on September 26-29th, 2001. The workshop was the fourth in a series of international conferences devoted to autonomous agents and multi-agent systems organized in Central-Eastern Europe. Its predecessors were CEEMAS'99 and DAIMAS'97, which took place in St. Petersburg, Russia, as well as DIMAS'95, which took place in Cracow, Poland.

Organizers of all these events made efforts to make them wide-open to participants from all over the world. This would have been impossible without some help from friendly centers in the Czech Republic, England, France, Japan, and The Netherlands. A total of 61 papers were submitted to CEEMAS 2001 from 17 countries. Out of these papers, 31 were selected for regular presentation, while 14 were qualified as posters. There were around 70 participants enjoying both the conference and the city.

The motto of the meeting was "Diversity is the core of multi-agent systems". This variety of subjects was clearly visible in the CEEMAS 2001 program, which addressed the following major areas of multi-agent systems:

- Organizations and social aspects of multi-agent systems
- Agent and multi-agent system architectures, models, and formalisms
- Communication languages, protocols, and negotiation
- Applications of multi-agent systems
- Agent and multi-agent development tools
- Theoretical foundations of Distributed AI
- Learning in multi-agent systems

The richness of workshop subjects was ensured thanks to the CEEMAS 2001 contributing authors as well as the keynote speakers from all over the world. They gave five interesting and highly diverse lectures announcing new trends in multi-agent technology:

- Jeffrey M. Bradshaw (University of West Florida, USA): Terraforming Cyberspace
- Toru Ishida (Kyoto University, Japan): Social Agents in Digital Cities
- Nick Jennings (University of Southampton, UK): Automated Haggling: Building Artificial Negotiators
- Andrzej Skowron (Warsaw University, Poland): Approximate Reasoning by Agents
- Krzysztof Zieliński (University of Metallurgy and Mining, Poland): Network Services in Context of Pervasive Mobile Internet.

The event was arranged and organized by a team of researchers from the University of Mining and Metallurgy and Jagiellonian University, led by Edward Nawarecki (general chair), supported by the other chairpersons: Barbara Dunin-Kêplicz from Warsaw University (program co-chair), Yves Demazeau from MAGMA, Grenoble (co-chair), and Vladimir Gorodetski from St. Petersburg Institute of Informatics and Automation (co-chair).

Recently, the CEEMAS'01 postproceedings "From Theory to Practice in Multi-Agent Systems" were published as Volume 2296 of Springer's Lecture Notes in Artificial Intelligence, under the editorship of Barbara Dunin-Keplicz and Edward Nawarecki. Taking advantage of the conference discussions, the authors improved their contributions, linking them with other CEEMAS'01 papers.

We would like to acknowledge the role of the workshop sponsors: the Rectors of the University of Mining and Metallurgy and the Jagiellonian University, as well as AgentLink. The AgentLink sponsorship was of a special importance: first of all, many Central-Eastern European scientific institutions took the opportunity to join AgentLink, now that it had become possible. Secondly, many authors from these countries submitted papers, being offered traveling grants. Those who usually have trouble finding resources to attend international conferences, know how precious such an opportunity is.

CEEMAS'01 was a very lively event: during many conference and evening discussions it was confirmed that the conference as an event was a great success, both scientifically and culturally. Scientifically, thanks to the efforts of outstanding invited speakers, contributing authors, the participants and the conference organizers. Culturally, this was rather easy as Cracow is not only the most beautiful city in Poland, but is also among the most charismatic places in Europe.

I am deeply convinced that the good atmosphere will stay with us during the next CEEMAS planned for Prague, Czech Republic, June 18-21, 2003. See: **http://cyber.felk.cvut.cz/ceemas2003/** for details of CEEMAS 2003.

The Second Symposium on Adaptive Agents and Multi-Agent Systems (AAMAS-II)

Eduardo Alonso
City University
UK
eduardo@soi.city.ac.uk

Daniel Kudenko
University of York
UK
kudenko@cs.york.ac.uk

Dimitar Kazakov
University of York
UK
kazakov@cs.york.ac.uk

The Second Symposium on Adaptive Agents and Multi-Agent Systems (AAMAS-II) was held on the 4th and 5th of April 2002 at Imperial College, London, as part of the AISB'02 convention. AAMAS-II was a continuation of AAMAS-I, held as part of AISB-01 in York, March 2001. AAMAS-I was a pioneering experience, as no symposium on learning agents had been organised previously in the UK. The success of the symposium encouraged us to organise AAMAS-II. Over 100 participants from Europe, America, Asia, and Australia took part in AAMAS-II. Sixteen contributions (eleven full papers and five short presentations) out of twenty-two submissions were accepted for presentation.

Sašo Džeroski, from Institut Jožef Stefan, Slovenia, gave an invited talk entitled "Relational Reinforcement Learning for Agents in Worlds and Objects". He was funded by AgentLinkII, as was AAMAS-I's invited speaker, Enric Plaza. The symposium focused on the following areas:

- Learning and adaptation in multi-agent systems: The ability to learn is especially important for an agent when there are other agents acting in the environment. An important open question is whether and how single-agent learning techniques can be adapted to, and applied in, a multi-agent setting.
- Logic-based learning: The ability to incorporate background knowledge into agent decision-making and learning processes is arguably essential for effective performance in complex, dynamic domains. Logic-based learning mechanisms such as explanation-based learning and inductive logic programming are being used to test this hypothesis.
- Learning and communication: When several learning agents work in a team it may be beneficial for them to co-operate not just on task achievement but also on the learning process itself. Clearly, communication is an important tool for such co-operation.
- Natural selection, language and learning: These three issues are inter-linked through evolutionary search for the best language bias used for learning.
- Evolutionary agents and emergent multi-agent structures: Genetic algorithms are a particular machine learning approach that has been successfully applied to social simulation and other multi-agent domains. Specific techniques are still under development. One focus of this research area is on observing emergent behaviours.
- Industrial applications of learning agents: Agent technology is already having a strong impact on various applications, including e-commerce, entertainment, human-computer interfaces, and industrial plant control. Many of these applications are being equipped with machine learning technology.
- Distributed learning: The major question in this area is how agents can learn in a collaborative way as a group. This is in contrast to the alternative view on multi-agent learning where agents in a group learn individually and separate theories are obtained.

With the AISB second symposium on adaptive agents and multi-agent systems the following goals were achieved:

- Increase awareness and interest in adaptive agent research in the European AI community and encourage further research: Even though Europe is the home to some of the strongest research groups in agent and multi-agent systems, research on adaptive agents has only recently started to receive increased attention.
- Encourage collaboration between ML and Agent system communities: There are many researchers in Europe that have a strong machine learning background but only little or no knowledge of agents, and vice versa. AAMAS-I and AAMAS-II presented an ideal opportunity for these two groups to meet and discuss potential collaborations.
- Give a representative overview of current research in the area of adaptive agents in Europe and world-wide.

We sincerely thank our invited speaker, Sašo Džeroski, and our programme committee, Kurt Driessens, Peter Edwards, Eugenio Oliveria, Michael Schroeder, Kostas Stathis, and Niek Wijngaards, and other reviewers, Chris Child, David Mobach, Ana Paula Rocha, and Sander van Splunter. We would like to thank the AgentLinkII network for their financial support.

NETTAB 2002 workshop on Agents in Bioinformatics

Paolo Romano
National Cancer Research Institute
Italy
paolo@ist.unige.it

Emanuela Merelli
University of Camerino
Italy
emanuela.merelli@unicam.it

1. Introduction

The NETTAB 2002 workshop on "Agents in Bioinformatics" was the second in a series of workshops [11] focused on the most promising and innovative Information and teleCommunication Technologies (ICT) tools and on their usability in bioinformatics.

NETTAB 2001[8] was devoted to a variety of new standards, for distributing information in a standardized way, such as the eXtensible Markup Language (XML)[13] and its extensions which are specialized in the biomedical field, and for remotely accessing network information systems, such as CORBA[6].

These tools were seen as a way of improving remote access to molecular biology databanks. The role of ontologies was especially emphasized as a mean for allowing consistent definitions of, and access to, the huge quantity of information that is currently managed by sequence, genome, proteome and other biologically oriented databases.

It is becoming increasingly clear that significant improvements can be achieved in the bioinformatics field by designing and implementing new ICT tools that are able to distribute the computation burden, while reducing the volume of data transferred. From this point of view, it is often felt that software agents can play a major role.

2. Possible roles of agents in bioinformatics

As is widely known, agents are autonomous problem-solving entities that can interact with others and respond to changing circumstances. In the last five years, the application of agent-based systems in different fields, such as e-commerce, transports and health, has been tested and its efficacy been demonstrated. However, little has been done in the life sciences domain.

Nonetheless, it is felt that in the Semantic Web, the integration of biological resources (data and their elaboration) will almost certainly require an understanding of agents and software robots. This is the case even if we agree that the real problem facing data integration is not related to the technology, but is instead concerned with getting everyone to agree on the meaning of the terms they use, so as to lead to the definition of common languages to describe information resources.

It has been demonstrated that agent based systems and applications can be designed to support bioscientists during the process of genome analysis and annotation [4][5][7]. They can effectively manage and improve the processes that are involved with protein structure prediction [9][12] and help integrate DNA microarray results [1].

NETTAB 2002 was a good opportunity to report on the results achieved in this area, to discuss the benefits (and drawbacks) that agent-based systems may bring to the bioinformatics domain, and also to provide a list of the research topics that should be tackled in the near future to make the deployment of bioinformatics agent-based systems a reality.

The workshop also aimed to:

- inspire collaborations between the communities of bioinformatics scientists and artificial intelligence experts;
- strengthen the relationships within bioinformatics community; and
- collect the latest ideas, achievements and proposals in the application of agents and multi-agent systems to the bioinformatics area.

3. Structure of the scientific programme

Since 2001, NETTAB workshops have been organized according to a well balanced formula that aims to:

- focus on a well defined, emerging, promising ICT technology that can efficiently support bioinformatics and the analysis of biological information;
- introduce the basic knowledge related to the technology under analysis;
- sketch the most promising features of the technology in bioinformatics domains;
- show some valuable examples and case studies;
- allow for as much discussion among participants as possible; and
- demonstrate how the technology works and how easily it can be implemented in bioinformatics service environments.

The NETTAB 2002 workshop was therefore based on invited lectures, oral communications, posters and tutorials.

The opening lecture by Ian Horrocks, University of Manchester, was devoted to the implementation of the Semantic Web and the related

role of agents. Invited talks were divided into three sessions. The first session set a common ground by introducing biologists and bioinformaticians to agents and multi-agent systems, while Paolo Ciancarini of the University of Bologna, and V. S. Subrahmanian of the University of Maryland, presented the basic aspects of the technology.

The second session was devoted to the analysis of possible roles of software agents in bioinformatics. It included a presentation by Luc Moreau of the University of Southampton on the MyGrid project[10], where agents will be widely employed in the setting up of an e-scientists workbench. It was interesting to see how agents were selected for their ability to cope with the dynamic and heterogeneous nature of the bioinformatics domain. In MyGrid, agents are to be used to cope with such tasks as personalization (choice of preferred and most trusted services, most frequent workflows), communication (support of multiple existing languages, standards and protocols, agents seen as Web Services) and negotiation (e.g., for Quality of Services).

Other presentations were given by Midori Harris of the European Bioinformatics Institute, who introduced the Gene Ontology and discussed the use of ontologies as a mean for enabling conceptual communication between agent based systems, and by Martin Bishop of the Human Genome Mapping Programme Resource Center, who introduced the HGMP-RC users service environment and discussed the needs of end users, thus pointing out again that the personalization of the work environment is a main issue in the network services.

The third session was devoted to some real applications and case studies in the field. Kevin Bryson of the Institut National Recherche Agronomique, presented Agmial, which is an extension of the well known GeneWeaver system. Agmial, which is currently under development, tries to overcome some limitations of its ancestor by adding standardized communication and security tools, such as Web Services and SSL encryption.

Keith Decker of the University of Delaware, presented BioMAS[3], a multi-agent system for genomic annotation. BioMAS is built upon DECAF, a toolkit for creation of multi-agent systems by the same author. It is based on four overlapping multi-agent organizations that share a common local KnowledgeBase and which are devoted to different elaborations: EST processing for gene determination, basic annotation, user querying and functional annotation for the determination of gene functions. BioMAS is available online.

Emanuela Merelli of the University of Camerino presented BioAgent, a prototype of a mobile agent system developed to completely decentralize the local task processing that presently characterizes a workflow of an in silico experiment, and free bioscientists from the need to continuously interact with remote services. BioAgent has a four-layered software architecture, consisting of core, service agents, bioagents and workflow layers. Each layer can be compositionally configurable so as to tailor the system for any of the existing biodata and bioservices providers. Security issues are avoided by using a hierarchical communication model that any active agent in the system must adopt. BioAgent is an open source project and will soon be available online www.bioagent.net [2].

Steffen Möller of the University of Rostock, the author of EDITtoTrEMBL[9], discussed the possible use of agent technology for the analysis of gene expression data in microarray studies and for the resolution of conflicts in protein studies.

Discussion among participants was also stimulated by a panel on the real applicability of agents and multi-agent systems to bioinformatics, while some technical aspects of the implementation of agents based systems were presented in three tutorials given by Michael Luck, Michael Schroeder and Andrea Omicini.

4. Acknowledgments

We wish to thank the Programme Committee for its useful suggestions and comments and for its support in the selection of oral communications and posters, and the Organizing Committee for its support in the organization of the workshop.

5. Further information

Further information about NETTAB workshops, including the full programmes and some presentations, can be found at the NETTAB website (**http://www.nettab.org/**).
The printed version of the abstracts book is sold-out, but electronic copies of papers are available: email the chairman at paolo@ist.unige.it.

6. References

[1] Angeletti, M., Culmone, R., Merelli, E., An intelligent agents architecture for DNA-microarray data integration, In Proc. NETTAB workshop on "CORBA and XML: Towards a bioinformatics integrated network environment", Genova, May 2001.
[2] BioAgent: **http://www.bioagent.net/**
[3] BioMAS: **http://udgenome.ags.udel.edu/herpes/**
[4] Bryson, K., Luck, M., Joy, M., Jones, D.T., Agent interaction for bioinformatics data management, Applied Artificial Intelligence, 15(10):917-947, 2001.
[5] Bryson, K., Luck, M., Joy, M., Jones, D.T., Applying agents to bioinformatics in GeneWeaver, in Proc. Fourth Intl Workshop on Collaborative Information Agents IV, Lecture Notes in Artificial Intelligence, 1860, pp. 60-71, Springer-Verlag (2000)
[6] CORBA: **http://www.corba.org/.**
[7] Decker, K., Zheng, X., Schmidt, C., A Multi-Agents System for Automated Genetic Annotation, In Proc. Fifth Intl Conf. On Autonomous Agents, pp. 433-440, Montreal, June 2001.
[8] Martin, A.C.R. Can we integrate bioinformatics data on the Internet? Trends in Biotechnology, 19 (9), 327-328.
[9] Möller, S., Leser, U., Fleischmann, W., Apweiler, R., EDITtoTrEMBL: A distributed approach to high-quality automated protein sequence annotation. Bioinformatics, 15(3):219-227, 1999.
[10] MyGrid: **http://www.mygrid.org.uk/.**
[11] NETTAB workshops: **http://www.nettab.org/.**
[12] Schroeder, M., Data explosion in biology: Agents to the rescue,
http://www.soi.city.ac.uk/~msch/talks/ice/sld001.htm
[13] **XML: http://www.w3.org/XML/.**

Fifth International Workshop on Cooperative Information Agents (CIA 2001)

Matthias Klusch
DFKI
Germany
klusch@dfki.de

Franco Zambonelli
University of Modena
Italy
franco.zambonelli@unimo.it

Onn Shehory
IBM
Israel
onn@il.ibm.com

Information agent technology has become one of the key technologies for the Internet and the World Wide Web. It emerged mainly as a response to the challenges of cyberspace from both the technological and human user perspective. Development of information agents requires expertise from different research disciplines such as Artificial Intelligence (AI), advanced databases and knowledge based systems, distributed information systems, information retrieval, and human computer interaction (HCI). Cooperative information agent research and development focused initially on accessing multiple, heterogeneous, and distributed information sources. Gaining access to these systems, through Internet search engines, application program interfaces, wrappers, and web-based screens has been an important focus of cooperative intelligent agents. Research has also focused on the integration of this information into a coherent model that combines data and knowledge from the multiple sources. Finally, this information is disseminated to a wide audience, giving rise to issues such as data quality, information pedigree, source reliability, information security, personal privacy, and information value.

The fifth event of the international workshop series on Cooperative Information Agents was held in Modena, Italy, September 6-8, 2001. It continued a tradition of capturing the intrinsically interdisciplinary nature of the above research area, by calling for contributions from different research communities and by promoting open and informative discussions on all related topics. As in previous years, the workshop featured a sequence of regular and invited talks given by leading experts in the field. In 2001, the topics of the talks were mainly on the challenges of ubiquitous and pervasive computing for information agents. These challenges are due in particular to the necessity of an efficient utilization, evolution, and trust management of information agents for user-oriented information search, provision, and visualization in networked computing environments with small, mobile, and embedded devices. A different issue concerns the potential of agent-based support of massive distributed data warehousing worldwide.

As a novelty, the CIA 2001 workshop issued two awards for best paper and best system innovation to acknowledge particularly significant advances in research and development respectively, in the area of information agents. Regarding the "CIA 2001 System Innovation Award" each nominee, as signaled by the program committee during the review process, was requested to demonstrate a running prototype at the workshop to the jury of the award. Both awards were given at the end of the workshop.

The first and second winners of the best system award were, respectively, "LEAP - Enabling FIPA agents on small devices" developed by Federico Bergenti et al. (University of Parma, Italy) and an efficient system for reliable agent communication in wireless environments developed by Heikki Helin et al. (Sonera Corporation, Finland). Each paper submitted to CIA 2001 was considered for nomination for the "CIA 2001 Best Paper Award" by the program committee. This year the paper "Optimality and Risk in Purchase at Multiple Auctions" by Onn Shehory (IBM Research, Israel) won the award.

Based on the success of former workshops and the fact that the workshop is now recognized as a well-established event in the area of agent-based computing, this year's organizers decided to hold the CIA workshop as a stand-alone event for the second time in its history. In past years, the event was usualy co-located with other major conferences in the domains of agent technology (Autonomous Agents 1998, ICMAS 2000), and Artificial Intelligence (IJCAI 1999). The CIA 2001 event was held in collaboration with the 28th conference on Very Large Datebases (VLDB) to continue providing a bridge between the areas of research in databases and intelligent agents.

CIA 2001 featured 5 invited, 13 regular, and 12 short papers selected from 63 submissions. This represented a 15% increase in the number of submissions compared with the last year's workshop which confirms the success of the CIA workshop and its force of attraction, independent of any co-location. The result of the peer-review of all contributions is included in this volume, and is rich in interesting, inspiring, and advanced work in research and development of intelligent information agents worldwide. All workshop proceedings have been published by Springer Verlag as Lecture Notes in Artificial Intelligence volumes 1202 (1997), 1435 (1998), 1652 (1999), and 1860 (2000), respectively. The sponsors of the CIA 2001 included NASA Goddard Space Flight Center, USA; NOKIA Research Center, USA; Swiss Life AG, Switzerland; AgentLink II, European Network of Excellence for Agent-Based Computing; AutoDesk Inc., USA; OTConsulting, Italy; and IFMAS, International Foundation for Multi-Agent Systems.

AgentLink supported this workshop by the provision of a limited set of student travel grants to enable PhD students from AgentLink II member nodes to present their papers at the workshop.
The presentations given at the workshop were structured into following eight topical sections.

- Information Agents for Mobile and Wireless Environments: Practical Issues and Directions
- Personal Assistance: Interaction and Avatars
- Information Search and Recommendation
- Data Warehousing and Mining
- Collaborative Information Agents: Systems and Applications
- Trading Internet Agents (1): Auctions
- Trading Internet Agents (2): Strategies, Negotiation, and Design
- Issues of Collaboration and Coordination

Also in keeping with its tradition this year's event featured four invited talks:

- Information Agents: The Social Nature of Information and the Role of Trust by Cristiano Castelfranchi (Italy)
- Information Agents for Mobile and Embedded Devices by Timothy Finin (USA)
- Interactive Integration of Information Agents on the Web by Yasuhiko Kitamura (Japan)
- Data Warehouse Quality and Agent Technology by Matthias Jarke (Germany)

The social program of the CIA 2001 workshop began with a reception at the Faculty of Engineering, where drinks and appetizers were served, and the participants were welcomed by the local chair and the dean of the university of Modena & Reggio Emilia. Participants were also offered a guided visit to Galleria Ferrari in Maranello, a visit to "Acetaia" in Vignola with balsamic vinegar offerings, and enjoyed a social dinner at the castle of the city of Vignola.

EASSS 2002 photos

Franco Zambonelli takes time out from organising AAMAS to contribute to the summer school.

John-Jules Meyer and Nigel Gilbert during a coffee break.

The EASSS '02 Social event.

The Fourth European Agent Systems Summer School (EASSS 2002)

Steve Munroe
University of Southampton
United Kingdom
sjm01r@ecs.soton.ac.uk

8-12 July 2002
Bologna
Italy

Eileen shares a joke with Wiebe van der Hoek and Yves Demazeau to the cool sounds of Italian jazz.

Henry Hexmoor enjoys a traditional Italian dinner with students.

Stefan Poslad makes a point to Tim Finin.

The fourth AgentLink European Summer School on Agent Systems (EASSS 2002) took place in Bologna, Italy, from July 8th to 12th on the campus of the Computer Science Department of the University of Bologna. The event was co-located with the First International Joint Conference on Autonomous Agents and Multi-Agent Systems, which took place the following week.

Narrowly avoiding the unseasonable downpours that marked AAMAS, the Summer School enjoyed sweltering conditions. Students enjoyed a full complement of 17 courses during the week-long summer school, ranging from introductory courses to more advanced courses covering specialist areas and applications. Each course was presented by leading experts in the field from both Europe and the US, and the quality of courses was excellent. Despite the heat, the lack of air-conditioning, and the construction taking place outside, it is fair to say that the School proved a tremendous success once again. Although feedback indicated that the air-conditioning was the main difficulty faced by students, Tuomas Sandholm proved what was possible in his course with a marathon stint of eight hours on the last day. The course continued until the very last minute of the day, and was thoroughly appreciated. Indeed, the quality of courses, and the efforts of all the lecturers, were hallmarks of the event, as always.

Aside from the technical programme, the highlight of the week was the social event at Cantina Bentivoglio, where participants were treated to a traditional Bolognese dinner, accompanied by a jazz band. For those who made it to the end, the evening was completed by glasses of grappa all round. But the surroundings of Bologna more generally provided a perfect environment for work and relaxation combined, throughout the week.

EASSS 2002 was jointly chaired by Paolo Ciancarini (University of Bologna), Wiebe van der Hoek (Universiteit Utrecht) and Michael Luck (University of Southampton, UK). The organizing committee also included Andrea Omicini (University of Bologna) and Franco Zambonelli (University of Modena), as well as Michael Wooldridge (University of Liverpool) and Gerhard Weiss (Technical University of Munich). There were 161 participants (including speakers) from across a range of 24 countries.

EASSS 2003 will be held in Barcelona from the 10th until the 14th of February. For details, see **http://www.agentlink.org/happenings/easss/2003/index.html.**

The Fourth Workshop of the UK Special Interest Group on Multi-Agent Systems (UKMAS)

December 2001

Mark d'Inverno
University of Westminster
London
UK
dinverm@wmin.ac.uk

The UKMAS workshops are concerned both with the dissemination of recent research within multi-agent systems and the provision of an appropriate forum for debate and discussion between students, academics and industrialists alike. As with UKMAS 2000, this year's event attracted around 80 participants and the format of this meeting was typical of previous years (Luck, 1997; Luck et al., 1998; Decker et al., 1999; Rana et al., 2000; d'Inverno et al., 2001).

Aside from the paper sessions, which spanned the entire range of work being undertaken acroos Europe, as well as the UK, the workshop included two invited talks. First, Samson Abramsky of Oxford University gave a talk entitled "From Computation to Interaction". As Information Technology has progressed from batch processing through multi-tasking operating systems to distributed systems and on to today's Internet, the focus has correspondingly shifted, from stand-alone programs to systems, and from computation to interaction. The first-generation models of computation gave an account of stand-alone programs computing functions from inputs to outputs. These are now seen as (very) special cases of a wider class of behaviours, in which the components of a complex system of concurrently executing agents interact with each other to achieve some global effect. On a macro scale, this description evidently fits today's distributed transactions across the Internet; but on a micro scale, it is equally true of how functional computations are ultimately realized, whether in software or in hardware. Ultimately, it seems that all computation can be resolved into the interactions of large numbers of very simple agents: the complex behaviour of the overall system is an emergent property of these interactions.

Second, Sarit Kraus gave a talk concerning Real Time Negotiation and Cooperation that covered the development of agents that can cooperate under time pressure. She described two projects, the first project concerned cooperative agents, and the second related to self-interested agents that could benefit from cooperation. First, the Distributed Dispatcher Manager (DDM), a system for managing large collections of dynamically changing tasks was described. In this work, tasks are distributed over large geographic areas and teams consist of very large groups of mobile and cooperative agents which have direct access to only local information about their immediate environment. The DDM's contributions include: realtime processes for combining partial results to form an accurate global solution; increased system fault tolerance; and scalability to very large task and agent problem domains. The second project described the development of an automated agent that can negotiate efficiently with humans. In this work, the environment is characterised by two negotiators, time constraints, deadlines, full information, and the possibility of opting out. The agent can play either role, with communications between agent and human taking place using a semi-formal language. The model used in constructing the agent is based on a formal analysis of the scenario using game theoretic methods and heuristics for argumentation. The agent receives messages sent by humans, analyzes them and responds.

Finally, there was also a panel discussion on the use of agents in e-commerce. The panelists were Andrew Byde from Hewlett-Packard Labs, Bristol, Marc-Philippe Huget from the University of Liverpool, David Lamper from Oxford University and Martin Kollingbaum from the University of Aberdeen. All panelists were asked three broad questions: their favourite current example of an agent system in e-commerce (that clearly illustrates the value-added nature of the agentification); the most compelling application of or role for agents in future generation e-commerce systems; and the key impediments (technical and/or social) to the widespread adoption of agent-mediated electronic commerce.

Once again UKMAS 2001 provided an ideal opportunity for agent researchers, practitioners and students to come together in an informal and lively but structured environment in order to present, debate and discuss current issues in the development of multi-agent research and development within the UK.

To coincide with this fifth event in the series, a book containing the best papers presented at the five workshops to date has just been published as a volume in the Springer Lecture Notes in Artificial Intelligence series (d'Inverno et al., 2002). Finally, the 2002 workshop will take place in Liverpool on the 18th and 19th of December 2002. For details, see **http://www.csc.liv.ac.uk/~mjw/research/ukmas/**

M. d'Inverno and M. Luck, 2001. Multi-Agent Systems Research into the 21st Century, *The Knowledge Engineering Review*, 16(3):271-275.

M. d'Inverno, M. Luck, M. Fisher and C. Preist, *Foundations and Applications of Multi-Agent Systems*, Lecture Notes in Artificial Intelligence 2403, Springer-Verlag, 2002.

K. Decker, M. Fisher, M. Luck, M. Tennenholtz and UKMAS'98 Contributors, 1999. Continuing Research in Multi-Agent Systems, *The Knowledge Engineering Review*, 14(3):279-283.

M. Luck, 1997. Foundations of Multi-Agent Systems: Issues and Directions, *The Knowledge Engineering Review*, 12(3):307-308.

M. Luck, M. d'Inverno, M. Fisher and FoMAS'97 contributors, 1998. Foundations of Multi-Agent Systems: Techniques, Tools and Theory, *The Knowledge Engineering Review*, 13(3):297-302.

O. Rana, C. Preist and M. Luck, 2000. Progress in Multi-Agent Systems Research, *The Knowledge Engineering Review*, 15(3):293-301.

Agents that Learn, Adapt and Discover

Eduardo Alonso
City University
United Kingdom
eduardo@soi.city.ac.uk

Amsterdam December 2001

The first ALAD (Agents that Learn, Adapt and Discover) SIG meeting was held in Amsterdam, December 2001. The meeting opened with an introduction by ALAD's co-ordinator, Pete Edwards from University of Aberdeen. He explained how the new SIG had been considered and approved by the AgentLInk II committee in Prague, Summer 2001.

This kick-off meeting continued with a presentation from Daniel Kudenko of the University of York, who described various research activities in the area of learning agents and multi-agent systems that are on-going in the York Artificial Intelligence research group. The research projects are in the areas of distributed learning, learning of coordination, evolution, information retrieval, and intelligent interface agents. Eduardo Alonso of City University then described a proposal to build and implement hybrid learning systems for e-commerce. Hybrid systems can learn either on-line (for example, using reinforcement learning techniques) or off-line (using domain knowledge and sets of examples to induce general hypothesis). The goal of the proposed research is to identify business scenarios where different learning strategies are successful.

Perhaps more importantly in relation to the SIG's objectives, there was also a talk from Georgios Tselentis of MIT GmbH, who described the organizational structure and goals of EUNITE, and offered possible overlaps with the interests of AgentLink and ILPNet in the areas of learning and adaptive agents.

The meeting finished with a planning discussion, where it was decided to organise the SIG in a centralised fashion, and leave the eventual creation of sub-groups for a later date. The discussion also covered the question of industrial involvement, and it was noted that there is still a lot to be done to overcome the gap between academia and industry. A possible activity that would bring both interests together would be the definition of benchmark tasks that are based on real-world data from industry.

Bologna July 2002

The joint AMEC/ALAD (Agent Mediated Electronic Commerce/ Agents that Learn, Adapt and Discover) SIGs joint session was held July 12th 2002, in Bologna. With this joint session the following goals were achieved:

- Increase awareness and interest in adaptive agents for e-commerce in the European AI community, and encourage further research: Even though Europe is the home to some of the strongest research groups in both AMEC and ALAD, research on adaptive agents for electronic commerce has only recently started to receive increased attention.
- Encourage collaboration between AMEC and ALAD communities: There are many agent researchers in Europe that have a strong learning background but only little or no knowledge of electronic commerce, and vice versa. The AMEC/ALAD meeting presented an ideal opportunity for these two groups to meet and discuss potential collaborations.
- Give a representative overview of current research in the area of adaptive agents for electronic commerce in Europe.

It is worth noting that the session was a truly international event, as researchers from France, Belgium, United Kingdom, The Netherlands, Spain, Portugal, Italy, Israel and USA participated. It was organised as an open discussion forum around different talks:

First, Ed Durfee of the University of Michigan discussed automated agents that learn consumer preferences in an information economy. Current highly networked information technologies permit people and enterprises to access vast amounts of information, and allow producers of information to supply a continuously evolving population of consumers who want their information goods. Given that an information producer is unlikely to know exactly what consumers will want and how much they want it, and that tracking such consumer preferences over time will require ongoing calculations, Ed has been investigating the use of computational agents to automate the producers' processes of adjusting contents

and prices of information products in response to consumer preferences. He summarized various experiences he had in trying to model and automate various business scenarios, ranging from how a monopolist producer can efficiently learn about a consumer population so as to maximize aggregate profit, to how a pair of producers can establish niches within a consumer population, to how a wider variety of markets for information goods can emerge in a larger population of producers and consumers.

Next, Henrique Lopes Cardoso and Eugénio Oliveira of the University of Porto talked about marketing strategies in an electronic commerce framework. B2C electronic commerce is commonly viewed as a straightforward browsing and purchase process. Current search tools assisting shoppers limit the automatic comparison of products to the price of goods. Allowing merchants to successfully employ differentiation strategies online requires providing the web with semantic content that can represent merchants' offers and be automatically processed by user-delegated software agents. In their talk, the authors explored what they believe would be a positive evolution of agent-based electronic commerce, justifying the need for "second generation bots" that enable the utilization of different seller-side marketing strategies. They take a dynamic posted offering approach as opposed to a single negotiation-based approach. They presented their seller agent architecture in such a scenario, including market analysis skills for insightful offer generation, implementing marketing strategies according to its configurable degree of autonomy. Since normally it will operate on imperfect information, they also consider the application of online learning skills, based on the exploitation/exploration tradeoff, that enable the agent to acquire information about the market in order to generate more informed offers.

Jean-Pierre Georgé, Marie-Pierre Gleizes and Pierre Glize of Institut de Recherche en Informatique de Toulouse then described a theory to handle adaptive MAS (Multi-Agent Systems), ABROSE, and presented a case study in electronic commerce. Adaptive MAS represent the next generation of MAS because most future applications will need to take into account the dynamic and complex nature of their environment. Classically such a system is defined as a MAS that is able to change its behaviour to react to the evolution of its environment. The authors' work focuses on the design of adaptive systems; therefore, they developed a theory based on emergence and MAS. The global functionality emerges from interactions between agents that can change the nature of their interactions with others in response to a changing environment. The case study in e-commerce is based on the exchange of products between producers and consumers and consists in brokerage and negotiation phases. The ABROSE system is a help for the brokerage phase by providing to a customer a list of relevant content providers, and in allowing a targeted publicity of the content providers' new offers. Each user is represented by a transaction agent whose behaviour is based on beliefs regarding itself and other agents. These beliefs are automatically updated after the user has done a transaction enabling ABROSE to be adaptive.

Rodrigo Agerri and Eduardo Alonso of City University, London, presented their work on learning to recover intentions in electronic commerce. The authors take a pragmatic view of inter-agent communication in MAS. In their view, agents' speech acts are actions whose purpose is the achievement of a goal. This goal is identified as the speaker's intention, which is expressed by using a sentence. Thus, 'understanding' is the process of working out the communicative meaning of an utterance or, in other words, the process of recovering agents' intentions from their uttered sentences. The authors' research focuses on learning how to recover such intentions. In particular, they use background knowledge and previous communicative experience to induce hypothesis about what the speaker's intentions are. Electronic commerce is an important domain for the implementation of their proposal, because when negotiating, it is essential to learn to recover the traders' intentions faithfully. In particular, in B2B scenarios where huge amounts of goods and resources are involved, understanding the "right" intention of one's counterpart in terms of price, quantity, quality, warranty and delivery time becomes crucial. The ability to automatically recover the intentions involved in the negotiation process might help understand and enforce dubious contracts.

Finally, there was a talk from Niek Wijngaards of Vrije Universiteit Amsterdam, discussing work with Brazier and Overeinder on automated adaption of e-commerce agents.

Agent Events Calendar

For further details visit: www.agentlink.org/happenings/other-events.php

If you would like to add your event email publications@agentlink.org with the details.

2002			
SEAL 2002	Singapore	Nov 18-22	4th Asia-Pacific Conference on Simulated Evolution And Learning
WDH 2002	Sydney, Australia	Dec 9-13	1st Workshop on the Modelling of Dynamical Hierarchies in Alife
Artificial Life VIII	Sydney, Australia	Dec 9-13	Artificial Life VIII: The 8th International Conference on the Simulation and Synthesis of Living Systems
Agent Day 2002	Belfort, France	Dec 12-13	Agent Day 2002
Symposium	Liverpool, UK	Dec 16-17	Symposium on Logic in Games and Multiagent Systems
UKMAS-2002	Liverpool, UK	Dec 18-19	5th Workshop of the UK Special Interest Group on Multiagent Systems (UKMAS-2002)
2003			
HICSS 2003	Big Island, Hawaii, USA	Jan 6-9	Hawaii International Conference on System Sciences
HICSS 2003 minitrack on Mobile Software Agents	Big Island, Hawaii, USA	Jan 6-9	Mobile Software Agents and their Use in Industrial Applications Minitrack Part of the Software Technology Track Thirty-sixth Annual Hawaii International Conference on System Sciences
UI 2003	Miami, USA	Jan 12-15	International Conference on Intelligent User Interfaces 2003'
AWIC-2003	Vienna, Austria	Feb 12-14	International Conference on Computational Intelligence for Modelling Control and Automation
CIMCA-2003	Vienna, Austria	Feb 12-14	International Conference on Intelligent Agents Web Technologies and Internet Commerce
IAWTIC 2003	Vienna, Austria	Feb 12-14	International Conference on Intelligent Agents, Web Technology and Internet Commerce - IAWTIC\'2003
SOCO/ ISFI 2003	Tenerife, Spain	Feb 12-15	5th International ICSC Symposium on Soft Computing and Intelligent Systems for industry
SAC2003-CM	Melbourne, Florida, USA	March 9-12	Coordination Models, Languages and Applications: Special Track of the 18th ACM Symposium on Applied Computing (SAC 2003
SAC 2003	Melbourne, Florida, USA	March 9-12	Special Track on E-Commerce Technologies
AMKM-03	Palo Alto, USA	March 24-26	Agent-Mediated Knowledge Management (AMKM-03) at Stanford University
AAAI 2003	Palo Alto, USA	March 24-26	The 2003 AAAI Spring Symposium Series
CHI 2003	Fort Lauderdale, USA	April 5-10	The CHI 2003 Conference on Human Factors in Computing Systems
IMS 2003	Budapest, Hungary	April 6-8	7th IFAC Workshop on Intelligent Manufacturing Systems
Agent-Based Simulation 4	Montpellier, France	April 28-30	The 4th International Workshop on Agent-Based Simulation
ICEIS 2003	Angers, France	April 23-26	5th International Conference on Enterprise Information Systems
CASA 2003	New-Brunswick, New Jersey, USA	May 7-9	The 16th International Conference on Computer Animation and Social Agents
FLAIRS-2003	St. Augustine, Florida, USA	May 12-14	Special Track on Integrated Intelligent Systems
Workshop at IEEE CCGrid 2003'	Tokyo, Japan	May 12-15	Agent Based Cluster and Grid Computing at IEEE CCGrid 2003
ICMAT2003	Chongqing, China	May 29-31	The Second International Conference on Active Media Technology
SABIS	Colorado Springs, USA	June 4-6	Software Agents in Business Information Systems
II2	Zakopane, Poland	June 2-5	Intelligent Information Systems 2003
AI '2003	Halifax, Nova Scotia, Canada	June 11-13	The Sixteenth Canadian Conference on Artificial Intelligence
CEEMAS 2003	Prague, Czech Republic	June 16-18	The 3rd International/Central and Eastern European Conference on Multi-Agent Systems
IEA/AIE-2003	Loughborough, UK	June 23-26	The 16th International Conference on Industrial & Engineering Applications of Artificial Intelligence and Expert Systems
IASSE-2003	San Francisco, USA	July 7-9	12th International Conference on Intelligent and Adaptive Systems and Software Engineering
AAMAS 2003	Melbourne, Australia	July 14-18	Second International Joint Conference on Autonomous Agents and Multiagent Systems
IJCAI-03	Acapulco, Mexico	Aug 9-15	The Eighteenth International Joint Conference on Artificial Intelligence
IDA 2003	Berlin, Germany	Aug 28-30	The 5th International Symposium on Intelligent Data Analysis
KES 2003	Oxford, UK	Sept 3-5	Seventh International Conference on Knowledge-Based Intelligent Information & Engineering Systems
IDAACS 2003	Lviv, Ukraine	Sept 9-10	IEEE Second International Workshop on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS\'2003)
IVA 2003	Irsee, Germany	Sept 15-17	The 4th International Working Conference on InIntelligent Virtual Agents
CINC '03	Cary, North Carolina, USA	Sept 26-30	The Sixth International Conference on Computational Intelligence & Natural Computation held under the auspices of JCIS 2003
JCIS 2003	Cary, North Carolina, USA	Sept 26-30	7th Joint Conference on Information Sciences
KIMAS '03	Cambridge, MA, USA	Oct 1-3	Integration of Knowledge Intensive Multi-Agent Systems
UbiComp 2003	Seattle, USA	Oct 12-15	The Fifth International Conference on Ubiquitous Computing
IAT 2003	Beijing, China	Oct 13-17	International Conference on Intelligent Agent Technology (IAT 2003)
WI 2003	Beijing, China	Oct 13-17	IEEE/WIC International Conference on Web Intelligence (WI 2003)
ISWC 2003	Sanibel Island, Florida, USA	Oct 20-23	2nd International Semantic Web Conference
2004			
4th EIS 2004	Maribor, Slovenia	June 28-2 July	International ICSC Symposium on Engineering of Intelligent Systems
ECAI 2004	Valencia, Spain	Aug 23-27	European Conference on Artificial Intelligence
2005			
IJCAI-05	Edinburgh, UK	July 31- 5 Aug	Nineteenth International Joint Conference on Artificial Intelligence