ISSN 1465-3842

AgentLink News 3

Europe's ESPRIT-funded Network of Excellence for agent-based computing

July 1999 **www.agentlink.org**

Features

Coordination Models for Multi-agent Systems
Paolo Ciancarini, Andrea Omicini and Franco Zambonelli *3*

Grasshopper - the Agent platform
Michael Höft, Thomas Magedanz and Joachim Quantz *7*

Facilitating the Deployment of Intelligent Agents in the Application Development Mainstream
Steven H. Schoepke *10*

Project reports

Agent Technologies for Federated Information Systems in the Public Administration
Maurizio Panti, Aldo Dragoni, Luca Spalazzi, Claudia Diamantini and Salvatore Valenti *12*

Just-in-time Information
Chris Knowles *14*

Announcements

Turning the Global Infrastructure into a "Universal Information Ecosystem"
Leonardo Flores *17*

Agents 2000 *18*

ICMAS 2000 *19*

AgentLink/SIG Reports

What's Happening in AgentLink?
Mike Wooldridge *20*

Agent Mediated Electronic Commerce - Second Meeting
Carles Sierra *21*

Methodologies and Software Engineering for Agent Systems - Second Meeting
Jan Treur *22*

Agent-based Social Simulation
Rosaria Conte *23*

Intelligent Information Agents - Second Meeting
Matthias Klusch *25*

Multi-Agent Coordination and Control
John Perram *27*

Agent R&D in Europe
Mike Wooldridge *28*

Conference and Workshop Calendar *30*

In order to understand and manage large multi-agent systems, it is necessary to model the interaction between agents as well as the internal behaviour of each agent. In other words, we need an appropriate coordination model, which is a high-level interaction abstraction aimed at globally ruling the behaviour of the different agents of the system. In this issue, Paolo Ciancarini and his colleagues try to motivate this need and briefly sketch the basic concepts of a coordination model.

Although agent technology today potentially offers superior solutions to many practical problems and shows even greater promise for future development, it is still not clear that it will achieve widespread deployment by software engineers outside research labs. Steven Schoepke has identified some of the main obstacles that agent technology will meet on its way to mainstream acceptance and suggests strategies for avoiding them. One of the proposed strategies concerns the issue of standardisation which, in turn, relates to the overview of the Mobile Agent System Interoperability Facility (MASIF) standard proposed by the Object Management Group (OMG) that was included in the first issue of AgentLink News. This article is here followed up by a description of Grasshopper, the first available MASIF-compliant mobile agent development and runtime platform. In this issue you will also find two project reports on agent-based information management systems and some notes about what is happening in AgentLink.

This issue also includes reports from the SIG meetings held in London, UK, on April 21-22, 1999. An important activity of these meetings was the construction of a "technological roadmap" for each particular field of interest. These roadmaps, which will play a crucial role in developing the AgentLink roadmap necessary for the continuation proposal, should include basically four things: an assessment of the state of the art, a long-term vision for the field, the technology gaps that exist between this long-term vision and the state of the art, and finally the short, medium, and long term research and development issues needed to be addressed in order to close these gaps.

Paul Davidsson, Editor

Coordination Models for Multi-Agent Systems

Paolo Ciancarini[1], Andrea Omicini[2] and Franco Zambonelli[3]
[1] Dipartimento di Informatica, Università di Bologna, Italy
[2] Dipartimento di Elettronica Informatica e Sistemistica, Università di Bologna, Italy
[3] Dipartimento di Scienza dell'Ingegneria, Università di Modena e Reggio Emilia, Italy
[1]ciancarini@cs.unibo-it, [2]aomicini@deis.unibo.it, [3]franco.zambonelli@unimo.it

The attempts to determine which features a software entity should present to be defined an "agent" at least agree on the fact that an "agent" is an autonomous entity which proactively interacts with its surrounding environment. A multi-agent system is then built by composing several independent computational activities, which have to interact with each other and with their environment. Therefore, a comprehensive view of a multi-agent system should not rely only upon the analysis of the internal behaviour of each agent, due to their intrinsically interactive nature. In fact, these systems are likely to exhibit a complex global behaviour, emerging from the mutual interaction among components, that is hard to be described and managed when communication is considered from a single agent's viewpoint.

In this paper, we claim that the design and management of a multi-agent system may take advantage from the choice and exploitation of a *coordination model*, that is, a high-level interaction abstraction aimed at globally ruling the behaviour of the different system components. In order to illustrate the issues arising from the choice of a suitable coordination model for the design of multi-agent systems, we take Internet applications as a case study.

What is needed to build a multi-agent system?

Even though our answer is a coordination model, we will not start from the definition of what a coordination model is and show how good it can be. Instead, we prefer to drive the reader through a short guided tour of some issues arising in the design and management of multi-agent systems, to let her to autonomously appreciate the usefulness of what we call a coordination model.

From Agent Theories to Multi-Agent Theories

Many research efforts in the area of autonomous and intelligent agents have focused on intra-agent aspects, such as agent languages and architectures, which focus on the agent's internal structure. Instead, multi-agent systems should not be simply considered as a multitude of individuals. If a society of agent exists, a collective social behaviour is likely to emerge. Therefore, to fully understand multi-agent systems, theories for agent societies are needed. These theories should define what is the world that hosts the society, which laws rule the world, and which are the individuals that can populate it. In addition, if any intelligent global behaviour can emerge from a system, there should be a place where it should be found and monitored. In a multi-agent world, this intelligence cannot reside inside agents only, but it should be somehow spread among agents and the interaction space among them. That is, the world where agents live is not composed of agents only, but also of the abstractions supporting the interactions, as well as of the history of these interactions. Therefore, a multi-agent system should define not only the world where agents live, but also the media that permeate the space and enable agent interactions.

Limits of Agent Communication Languages

Agents usually interact by exchanging complex symbolic information and possibly have to agree on complex interaction protocols. In addition, agents are autonomous, possibly designed separately in different times by different people, and including heterogeneous software components. These issues have led to the development of several proposals in the area of agent communication languages [Sin98], such as the FIPA and the KQML proposals. From a different perspective, several middleware systems, notably CORBA, have been proposed as software layers enabling interoperability among software components. Although these works are important to achieve interoperability, they mainly focus on peer-to-peer communications and do not account for a more comprehensive view of the interaction as a primary component of agents' societies. Therefore, both agent communication languages and middleware systems have to somehow be extended in scope in order to include not only language and protocol specifications but also the definition of *coordination laws*, to allow for a global understanding and management of interactions.

Software Architectures for Multi-Agents Systems

When a multi-agent system is made up of a large number of independently designed components, is may be very difficult to correctly design and managing the system as a whole. An approach that simply puts components together and lets them to interact is likely to degenerate into chaos. Instead, models and tools are needed to put components together in a structured way: as already recognised in the area of software engineering, the design and management of a large software project requires the definition and analysis of its software architecture. This includes defining the role of each component, the mechanisms upon which composition can be based, and their composition laws. A similar approach would be undoubtedly helpful also in the context of multi-agent systems. However, in this case, a more dynamic and flexible definition of the software architecture - interaction-oriented rather than composition-oriented - is needed.

What is a coordination model?

In its general terms, coordination is the art of managing interactions and dependencies among activities, that is, in the context of multi-agent systems, among agents [Cia96, GelC92]. A coordination model provides a formal framework in which the interaction of software agents can be expressed. Generally speaking, a coordination model deals with the creation and destruction of agents, their communication activities, their distribution and mobility in space, as well as the synchronisation and distribution of their actions over time.

More precisely, a coordination model consists of three elements (Figure 1):

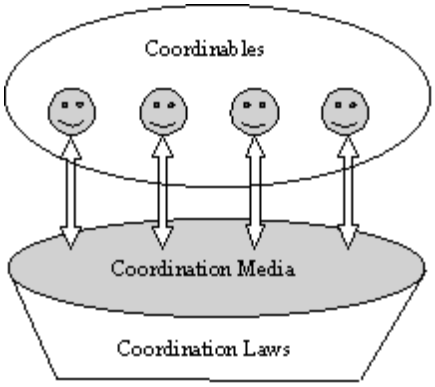

Figure 1. The components of a coordination model

- the *coordinables*, which are the entities whose mutual interaction is ruled by the model. These could be processes, threads, objects, even users and, as of interest here, agents;
- the *coordination media*, which are the abstractions enabling agent interactions, as well as the core around which the components of a coordinated system are organised. Examples are the classic media like semaphores, monitors, or channels, or more complex media like tuple spaces, blackboards, pipes, etc.
- the *coordination laws*, which define the behaviour of the coordination media in response to interaction events. The laws can be defined in terms of a communication language, that is a syntax used to express and exchange data structures, and a coordination language, that is a set of interaction primitives and their semantics.

From a software engineering viewpoint, a coordination model works as a source for design metaphors, abstractions, and mechanisms effectively supporting the definition of the software architecture and the development process of a multi-component software system.

Data-driven vs. control-driven coordination

An interesting survey and analysis of several coordination models, based upon the definition of a simple yet effective taxonomy, can be found in [PapA98]. Accordingly to that taxonomy, coordination models can be classified as control-driven or data driven ones.

In a *control-driven* coordination model (Figure 2-left), coordinables (agents) typically open themselves to the external world and interact with it through events occurring on well-defined input/output ports. The observable behaviour of the coordinables from the point of view of the coordination media is then the one of state changes and events occurring on these ports. For example, an event could be the announcement by a process to send out data from a given output port. The coordination laws establish how events and state changes can occur and how they should propagate. Therefore, the coordination media handle the topology of the interaction space among agents, without paying any attention to the data possibly exchanged between processes. From the viewpoint of the coordination media, the only data of interest are the communication events.

In *data-driven* coordination models (Figure 2-right), coordinables interact with the external world by exchanging data structures through the coordination media which basically act as shared data spaces. The coordination laws establish how data structures should be represented and how they should be stored and extracted from the data space. Unlike control-driven coordination models, the coordination media has no perception of the state changes of the coordinables and does not provide for any virtual connection among coordinables.

Different application contexts exhibit different needs with respect to coordination, and the choice of a coordination model is likely to have a great impact in the design of multi-component applications. In general, control-driven coordination models suits better those systems made up of a well-defined number of entities in which the flow of control and the dependencies between the components have to be regulated, and in which the data exchanged in not so important. These include, for example, computational intensive parallel applications, distributed management systems, definition of complex software architectures. Instead, data-driven model, seems to better suit open applications, where a number of possibly a-priori unknown and autonomous entities have to cooperate. In this case, focusing on dependencies between the components, as a control-driven model would do, would somehow clash with the autonomy of the components and the dynamicity of the open environment. Focusing on data preserve autonomy and dynamicity of autonomous components, which are usually designed to acquire information rather than control, as in the case of software agents.

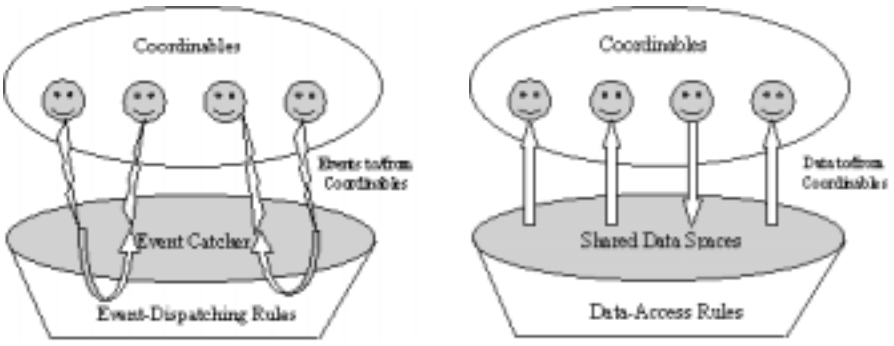

Figure 2. Control-driven (left) versus data-driven (right) coordination models

A data-driven coordination model

To better illustrate what a coordination model can be, in concrete, we recall the most notable example of a data-driven coordination model, namely the shared tuple space as in Linda. With regard to communication media, Linda relies on the notion of *tuple space* as an *associative blackboard* used by active components to communicate, synchronise, and cooperate by reading, writing, and consuming *tuples*. The communication language is then made up of tuples, which are simply ordered sets of typed fields. The coordination language provides coordinables with the primitives for tuple space access and modification. In particular, it includes an output operation, to write a tuple in a tuple space, and two input operations, to retrieve a tuple from the tuple space, respectively either deleting it or just copying it from the tuple space. Retrieving on tuples relies on an associative mechanism embedded into the tuple space itself: the input operations express their desiderata via a template tuple, and if a tuple corresponding to the template (eg. a *matching* tuple) is found in the tuple space, it is returned to the operation. Since the semantics of input operations imply that a coordinable agent blocks until a matching tuple is found, the Linda coordination language provides for synchronisation, too.

The tuple space metaphor is best understood as a coordination model embedded in a conventional programming language. For instance, when Linda is embedded in Java, coordinables are active Java objects, the coordination medium is a multiset of Java objects, the coordination laws are those describing the semantics of Linda-like primitives on reading, writing, and consuming Java objects in the tuple space.

After its first definition, many extensions, derivations and implementations have followed the original model, trying to exploit and maximise the benefit of its many features, which are nevertheless intrinsic in the model and have made Linda a widely known and exploited model. These include:

- *scalability* of the coordination media, possibly made up of a multiplicity of independent tuple spaces;
- *cleanness and expressiveness* of the communication language, which can provide three simple to understand yet powerful primitives;
- clean *separation of concerns* between computation and coordination activities, the formed transparently enclosed in the coordinables, the latter explicitly occurring though the tuple spaces;
- generative communication and associative access to communication information, leading to *spatial and temporal uncoupling* between coordinables: coordinables should not know each other to coordinate, neither they should synchronise.

Therefore, while Linda originates in the field of parallel programming, its underlying coordination model is well-suited also for open, distributed systems and, in particular, for complex multi-agent systems.

A Case Study: the Internet

For any coordination model to find an appropriate use in a given context, the above abstract issues must find a concrete design solution. This implies defining the coordination model and the language constructs for enabling its use. Since any discussion on this topic cannot abstract from the particular application environment, in this section we assume the Internet as a case study. Though it is now recognised that the design of interactive Internet applications can effectively exploit agents and Linda-like coordination [CTV99], even in the case agents are mobile [CabLZ99], several different choices are possible in the definition of the coordination architecture and language. This is confirmed by the different design choices adopted by three systems independently developed by the authors of this report, namely PageSpace [Cia98], TuCSoN [OmiZ99], and MARS [CabLZ98].

Coordinables

Apart from application specific mobile agents, an Internet application usually integrates several non-mobile entities, such as WWW servers and CORBA-compliant services, which an agent may be in need to interact with in order to perform its task. The problem is to determine the role these entities should play in the coordination model. A possible choice is to consider all these non-mobile entities as simply server objects, by defining their coordination medium as merely *another Internet service*, as in TuCSoN. As an opposite approach, the coordination architecture could be designed as to become *"The Internet" service*, masking and providing access to any existing service via the tuple space itself, as in MARS. Between the above two extremes, the approach adopted by PageSpace conceives the coordination model as a sort of *middleware*, selecting as coordinables all the entities populating the Web and offering a uniform underlying infrastructure. However, since legacy entities have not been explicitly designed to access a tuple space, their integration into the architecture may require special-purpose agents in charge of mediating and ruling the access to the tuple spaces from other entities.

Coordination Media

Building an Internet application by using a single tuple space is not feasible. The availability of independent tuple spaces associated to Internet sites enables decentralisation and modularity, but introduces novel problems. In particular, a coordination model including multiple tuple spaces requires for agents a suitable way to denote and access the spaces themselves.

A *transparent* - location independent - approach (adopted by PageSpace) make application agents refer to tuple spaces via references, transparently to location. An i*mplicit* - location dependent - approach (adopted by MARS) alleviates the agent from the need of explicitly referring to a tuple space, and automatically lets the agent access a default tuple space, possibly depending on its current position in the Internet. Therefore, the implicit approach tightly connects the issue of agent mobility to the coordination model. The *explicit* approach (adopted by TuCSoN), enables remote access to tuple spaces by associating a global - location aware - naming scheme to tuple space and it is very suitable for Internet applications, especially if the naming scheme is Web compliant one, such as URLs-based ones. However, when adopting an explicit approach, the implicit approach has not to be necessarily discarded: although it is always possible to refer to a tuple space by its global identifier, the possibility for an agent to implicitly refer to a "default" tuple space, depending on its current location, may reduce the complexity of the agent and permit applications to be based on "context-dependent" coordination actions (*explicit*).

Coordination Language

When adopting a coordination model based on tuple spaces, defining the coordination language means to define the *data types* and *language primitives* to access the spaces. With regard to the latter point, the original Linda model, mostly oriented to parallel programming, defines a simple language based on a limited set of typed primitives, to be used to compose tuples. The current trend, followed by PageSpace and MARS, is to define *object-oriented tuple space* models, in which both tuple spaces, as well as tuples and their component fields are Java objects. In some cases, however, other tuple space models may be preferable. For instance, since TuCSoN is oriented to the coordination of intelligent information agents, it defines a *logic tuple space* model in which tuples are terms, namely uninterpreded first-order logic predicates. Concerning the communication primitives, most of the systems define a static and *limited set of primitives* to be used to access the coordination media. Some systems enable agents to add any new primitives to a tuple space, in order to implement any needed operation on the tuples stored in it. However, this approach cannot generally suit Internet applications, based on autonomous agents, because it forces an agent to know which operation the other agents have installed in the tuple space.

Coordination Laws

In general, tuple spaces embed fixed, statically defined, and *not modifiable coordination laws*, that is, associative access to data, as in the case of PageSpace. However, due to the lack of flexibility of this approach, other proposals, including TuCSoN and MARS, conceive the coordination medium as a configurable kernel, by allowing *new coordination laws* to be defined and embedded in it [Den98]. In particular, in TuCSoN, coordination media can programmed in Prolog language by intelligent, meta-level agents, in order to dynamically tune the global behaviour of a coordinated application and embed any needed coordination laws into the tuple spaces. In its turn, MARS, coherently with its OO tuple space model, follows an object-oriented approach to coordination laws programming. The execution of one method of a class can be associated to specific communication event to override the default associative mechanism and implement specific behaviour of the tuple space in reaction to these kind of communication events.

Conclusions

The definition and exploitation of an appropriate coordination model is necessary to understand and manage large multi-agent systems. In this paper, we have tried to motivate this need and have briefly sketched the basic concepts for a coordination model. In particular, the paper should have made clear that a large number of different coordination models can be defined and, for each model, different architectural and programming language solutions exist to enable the use of the model in applications. The authors are currently involved in the analysis and implementation of Linda-like coordination models for Internet agents, which represents a challenging issue toward the effective exploitation of the Internet as a globally distributed computing systems.

References

[CabLZ98] G. Cabri, L. Leonardi, F. Zambonelli, Reactive Tuple Spaces for Mobile Agent Coordination, 2nd International Workshop on Mobile Agents, Lecture Notes in Computer Science, No. 1477, pp. 237-248, September 1998.

[CabLZ99] G. Cabri, L. Leonardi, F. Zambonelli, Coordination Models for Internet Applications based on Mobile Agents, IEEE Computer, 1999, to appear.

[Cia96] P. Ciancarini, Coordination Models and Languages as Software Integrators, ACM Computing Surveys, Vol. 28, No. 2, pp. 300-302, June 1996.

[Cia98] P. Ciancarini, R. Tolksdorf, F. Vitali, D. Rossi, A. Knoche., Coordinating Multi-Agents Applications on the WWW: a Reference Architecture, IEEE Transactions on Software Engineering, Vol. 24, No. 8, pp. 362-375, May 1998.

[CTV99] P. Ciancarini, R. Tolksdorf, F. Vitali, The World Wide Web as a Place for Agents, in M. Wooldridge and M. Veloso (eds.), Artificial Intelligence Today, Lecture Notes in Artificial Intelligence 1600 July, 1999 (to appear).

[Den98] E. Denti, A. Natali, A. Omicini, On the Expressive Power of a Language for Programming Coordination Media, 1998 ACM Symposium on Applied Computing, February 1998.

[GelC92] D. Gelernter, N. Carriero, Coordination Languages and Their Significance, Communications of the ACM, 35:2(96-107), February 1992.

[OmiZ99] A. Omicini, F. Zambonelli, Coordination for Internet Application Development, Journal of Autonomous Agents and Multi-Agent Systems, 1999, to appear.

[PapA98] G. A. Papadopoulos and F. Arbab, Coordination Models and Languages, Advances in Computers, Vol. 46, pp. 329-400, Academic Press, Aug. 1998.

[Sin98] M. P. Singh, Agent Communication Languages: Re-thinking the Principles, IEEE Computer, 31:12(40-47), December 1998.

Grasshopper - the Agent Platform

Michael Höft, Thomas Magedanz and Joachim Quantz
IKV++ GmbH, Germany
{hoeft, magedanz, quantz}@ikv.de

In the last few years, research in agent technology has addressed two issues which are of special importance for commercial exploitation, namely standardisation and the relationship between agent technology and object oriented technology. Today there is a common understanding that agent technology should be seen as an enhancement of distributed object technology and first agent standards are available. In particular, the *Object Management Group (OMG)* has produced the MASIF standard: Mobile Agent System Interoperability Facility [1] (see also the MASIF Overview in *AgentLink News 1*). Together with the more recent work within the newly established OMG Agent Special Interest Group this can be regarded as a major milestone on the road towards a unified distributed intelligent mobile object middleware. Another important standardisation forum for agent technologies and applications is the *Foundation of Intelligent Physical Agents (FIPA)* [2].

This article provides a brief overview of Grasshopper, which is the first available MASIF-conformant mobile agent platform today [3]. The first commercial Grasshopper version has been released in summer 1998 and is used today in several European projects, such as ANIMA, CAMELEON, FACTS, MARINE, MARINER and MIAMI. Most of these projects belong to CLIMATE, the Cluster for Intelligent Mobile Agents in Telecommunication Environments, which is part of the European Research Programme ACTS (Advanced Communications Technologies and Services) [4] (see also the CLIMATE Overview in *AgentLink News 2*).

We also sketch extensions envisaged for future releases of Grasshopper, such as the integration of a FIPA-Add-On making Grasshopper FIPA-compliant and supporting the Agent Communication Language ACL. Given Grasshopper's current functionality and the scope of future releases, Grasshopper can thus be characterised is an open agent platform, which offers through its standard conformance and the related interfaces a maximum of flexibility for the development of distributed applications.

Technical Overview

Grasshopper is a mobile agent development and runtime platform that is built on top of a distributed processing environment. It is based on Java JDK 1.1 and written in Java. This way, an integration of the traditional client/server paradigm and mobile agent technology is achieved. Most importantly, Grasshopper is developed compliant to the OMG MASIF specifications. Grasshopper can be used in various application domains, such as telecommunications, active networking and electronic commerce [5, 6, 7].

Grasshopper realises a *Distributed Agent Environment* (DAE). The DAE is composed of regions, places, agencies and different types of agents. Fig 1 depicts an abstract view of these entities.

Two types of agents are distinguished in Grasshopper: mobile agents and stationary agents. The actual runtime environment for both mobile and stationary agents is an *agency*: on each host at least one agency has to run to support the execution of agents. A Grasshopper agency consists of two parts: the core agency and one or more places. Core Agencies represent the minimal functionality. The following services are provided by a Grasshopper core agency:

Communication Service

This service is responsible for all remote interactions that take place between the distributed components of Grasshopper, such as location-transparent inter-agent communication, agent transport, and the localization of agents by means of the region registry. All interactions can be performed via CORBA IIOP, Java RMI, or plain socket connections. Optionally, RMI and plain socket connections can be protected by means of the Secure Socket Layer (SSL) which is the de-facto standard Internet security protocol. The communication service supports synchronous and asynchronous communication, multicast communication, as well as dynamic method invocation. As an alternative to the communication service, Grasshopper can use its OMG MASIF-compliant CORBA interfaces for remote interactions. For this purpose, each agency provides the interface *MAFAgentSystem*, and the region registries provide the interface *MAFFinder*.

Registration Service

Each agency must be able to know about all agents and places currently hosted, on the one hand for external management purposes and on the other hand in order to deliver information about registered entities to hosted agents. Furthermore, the registration service of each agency is connected to the region registry which maintains information of agents, agencies and places in the scope of a whole region.

Management Service

The management services allow the monitoring and control of agents and places of an agency by (human) users. It is possible, among others, to create, remove, suspend and resume agents,

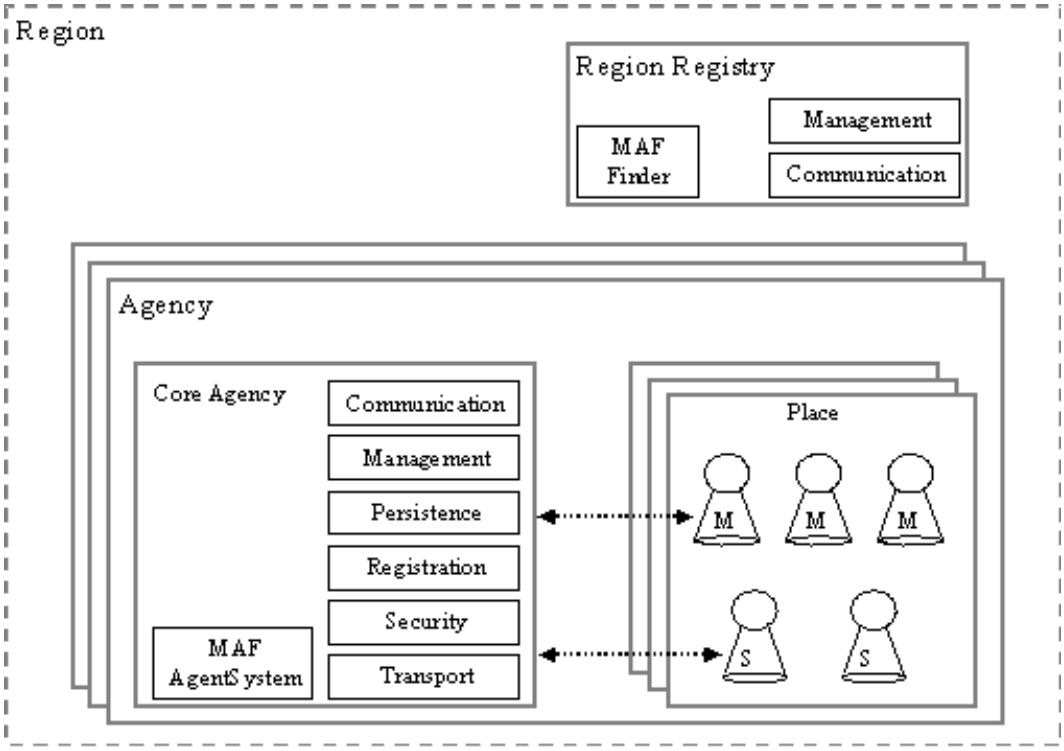

Fig. 1. Hierarchical Component Structure

services, and places, in order to get information about specific agents and services, to list all agents residing in a specific place, and to list all places of an agency.

Security Service

Grasshopper supports two security mechanisms: external and internal security.

- External security protects remote interactions between the distributed Grasshopper components, i.e. between agencies and region registries. For this purpose, X.509 certificates and the Secure Socket Layer (SSL) are used. SSL is an industry standard protocol that makes substantial use of both symmetric and asymmetric cryptography. By using SSL, confidentiality, data integrity, and mutual authentication of both communication partners can be achieved.
- Internal security protects agency resources from unauthorised access by agents. Besides, it is used to protect agents from each other. This is achieved by authenticating and authorising the user on whose behalf an agent is executed. Due to the authentication/ authorisation results, access control policies are activated. The internal security capabilities of Grasshopper are mainly based on JDK security mechanisms.

Persistence Service

The Grasshopper persistence service enables the storage of agents and places (the internal information maintained inside these components) on a persistent medium. This way, it is possible to recover agents or places when needed, e.g. when an agency is restarted after a system crash.

A place provides a logical grouping of functionality inside of an agency. The region concept facilitates the management of the distributed components (agencies, places, and agents) in the Grasshopper environment. Agencies as well as their places can be associated with a specific region by registering them within the accompanying region registry. All agents that are currently hosted by those agencies will also be automatically registered by the region registry. If an agent moves to another location, the corresponding registry information is automatically updated.

The functionality of Grasshopper is provided on the one hand by the platform itself, i.e. by *core agencies* and *region registries*, and on the other hand by *agents* that are running within the agencies, in this way enhancing the platform's functionality. The following possibilities regarding the access to the Grasshopper functionality must be distinguished:

- Agents can access the functionality of the *local agency*, i.e. the agency in which they are currently running, by invoking the methods of their super classes Service, StationaryAgent, and MobileAgent, respectively. These super classes are provided by the platform in order to build the bridge between individual agents and agencies, and each agent has to be derived from one of the classes StationaryAgent or MobileAgent.

- Agents as well as other DAE or non-DAE components, such as user applications, are able to access the functionality of *remote agencies* and *region registries*. For this purpose, each agency and region registry offers an external interface which can be accessed via the Grasshopper communication service.
- Agencies and region registries may optionally be accessed by means of the MASIF-compliant interfaces `MAFAgentSystem` and `MAFFinder`.

In the context of Grasshopper, each agent is regarded as a *service*, i.e. as a software component that offers functionality to other entities within the DAE. Each agent/service can be subdivided into a common and an individual part. The common (or core) part is represented by classes that are part of the Grasshopper platform, namely the classes `Service`, `MobileAgent`, and `StationaryAgent`, whereas the individual part has to be implemented by the agent programmer.

A Grasshopper agent consists of one or more Java classes. One of these classes builds the actual core of the agent and is referred to as *agent class*. Among others, this class has to implement the method live which specifies the actual task of the agent. The agent class must be derived either from the class StationaryAgent or from the class MobileAgent which in turn inherits from the common super class Service. The methods of these classes represent the essential interfaces between agents and their environment. The following two ways of method usage have to be distinguished:

- One part of the super class methods of an agent enable the access to the local core agency. For example, an agent may invoke the method `listMobileAgents()`, which it inherits from its super class Service, in order to retrieve a list of all other agents that are currently residing in the same agency.
- The remaining super class methods are defined to access individual agents. These methods are usually invoked by *other* agents or agencies via the *communication service* of Grasshopper. For instance, any agent may call the method `getState()` of *another* agent in order to retrieve information about the other agent's actual state. Note that this way of access is not performed directly on an agent instance, but instead on an agent's *proxy* object.

Two other ways of platform access are available for Grasshopper agents:

- The class RegionRegistrationP enables an agent to access the region registry in order to retrieve information about registered components, i.e. agencies, places, services, and agents. The region registry is accessed by means of a corresponding proxy object via the communication service.
- Apart from Grasshopper-specific platform access, the MASIF-compliant interfaces may be used, i.e. MAFAgentSystem for agencies and MAFFinder for the region registry.

Grasshopper Releases

Since February 1999, Grasshopper Release 1.2 is available, which could be mainly considered as a maintenance release providing performance optimisation as well as stability enhancements. Furthermore it offers a simplified user interface for the security configuration and the communication protocols are realised via a defined plug-in mechanism which enables the easy integration of additional communication protocols.

Besides the usual maintenance improvements, the upcoming Release 2.0 of Grasshopper will offer the following new features:

- complete support of the Java 2 Platform (formerly JDK 1.2)
- provision of an alternative mechanism as an optional substitute for the region registry, based on Jini/JNDI [8]
- Grasshopper Runtime environment as a Java applet
- provision of beans extensions for the major classes of the Grasshopper platform to enable the easy integration into beans supporting IDEs (Integrated Development Environments)

For autumn 1999 IKV++ also plans the release of a FIPA-Add-On for Grasshopper, which will allow the development of FIPA-compliant stationary agents. This Add-On will comprise the basic functionality required for a FIPA platform, namely an Agent Communication Channel (ACC), a Directory Facilitator (DF) and an Agent Management System (AMS). One of the main characteristics of the FIPA standard is the Agent Communication Language (ACL), a standardised format for messages sent between agents. To support these messages, the Add-On includes an ACL parser and additional tools for manipulating and visualising ACL messages. Furthermore, additional support for an XML format of ACL messages and for the XML encoding of message contents is envisaged.

Conclusion

This short article described the Grasshopper agent platform as an enabling technology for future service environments. Due to the envisaged openness of these environments (i.e., the high number of roles involved in service provision), the use of agent standards is fundamental. Grasshopper is an agent platform compliant to existing agent standards and thus provides a high degree of flexibility for software developers to realise their ideas. The Grasshopper agencies and places offer a powerful mechanism to structure complex software systems. Therefore Grasshopper is a strong candidate for the implementation of future agent based applications. Visit the Grasshopper web site [3] in order to obtain the free Grasshopper Light version. Furthermore, the web site provides related agent papers and technical documentation.

References

[1] OMG (1995), Common Facilities RFP3, Request for Proposal OMG TC Document 95-11-3, Nov. 1995, http://www.omg.org/; MASIF specification is available through http://ftp.omg.org/pub/docs/orbos/97- 10-05.pdf

[2] FIPA web site: http://www.fipa.org

[3] IKV++ GmbH – Grasshopper web page: http://www.ikv.de/products/grasshopper and http://www.ikv.de/download/

[4] ACTS CLIMATE web page: http://www.fokus.gmd.de/research/cc/ima/climate/climate.html

[5] M. Breugst et al. (1998). 'On the Usage of Standard Mobile Agent Platforms in Telecommunication Environments', Proceedings of 3rd ACTS Conference on Intelligence in Services and Networks (IS&N), Springer Verlag, ISBN: 3-540-64598-5, 1998, pp. 275 – 286

[6] M. Breugst, T. Magedanz (1998). 'Mobile Agents - Enabling Technology for Active Intelligent Network Implementation'; IEEE Network Magazine, August 1998, Vol. 12, No. 3, pp. 53-60,

[7] L. Hagen, T. Magedanz, M. Breugst (1998). 'Impacts of Mobile Agent Technology on Mobile Communications System Evolution'; IEEE Personal Communication Magazine; August 1998 Vol.5 No.4, pp.56-69

[8] http://www.javasoft.com/jdc/products/jini

Facilitating the Deployment of Intelligent Agents in the Application Development Mainstream

Steven H. Schoepke
Fannie Mae, Knowledge Engineering Group, Washington, D.C., USA
steven_h_schoepke@fanniemae.com

A lucrative window of opportunity currently exists for the field of Artificial Intelligence (AI) to emerge as a worthy force in the application development mainstream. Recent developments in academia and industry indicate a possible convergence of AI and the object-oriented/component-based paradigm. The common conduit to attain this objective could be the intelligent agent. However, one of the most daunting obstacles to successful deployment of this technology may very well be this movement's primary target: the developer and user populations. Their unfamiliarity with AI-related concepts and possibilities may present the greatest hindrance to widespread acceptance. The technical eloquence of intelligent agents may not be enough to achieve widespread deployment. In a "market" crowded with competitive solutions, any one offering cannot afford to rest on the laurels of its technical excellence, but must make optimal use of strategies on both the academic and the vendor sides of the "fence".

Many potential obstacles to AI's success could ultimately be overcome by the employment of prudent strategies, or a situation similar to the unmet expectations of the 1980s may result. During that period, the lack of an effective forum for communication and understanding, coupled with a lack of available vehicles (like intelligent agents) may have contributed to AI's failure to penetrate the mainstream. However, many recent initiatives, innovations, and implementations in the intelligent agent arena have resurrected AI's promising future. The following strategies attempt to realize this potential.

Standardization movements

An alliance (or at least a meeting of the minds) between FIPA (Foundation for Intelligent Physical Agents) and OMG (Object Management Group - the OO and UML standards organization) could result in a symbiotic fusion whereby strengths are complemented and weaknesses are compensated between the two groups. It's most efficient not to have to reinvent the wheel, and risk settling with a lesser version. A working and cooperative relationship between the groups to achieve some semblance of compromise on standards can be a critical first step in the advancement of intelligent agents and symbiosis between OO and AI. FIPA seems to have brought various AI disciplines into practical settings, and OMG has pioneered a well-established methodology supporting patterns and reuse. Right now their goals may not be mutually exclusive, and a cooperative collaboration could result in new synergies. As a metaphorical case in point, the human species could not have survived without common patterns of skeletons, organs, reasoning, and learning.

Alignment (and convergence) with the prevailing object-oriented/ component-based paradigm

A very fortunate connection exists between the AI-related technology embedded in intelligent agents and Java. The latter, being object-oriented in nature, could transfer to the AI components some very important constructs: reuse, design and analysis patterns, organized and disciplined methodologies, the power of inheritance, and maintenance/enhancement advantages.

An effective alignment effort would involve representing and integrating the various AI-related technologies into currently accepted models and methodologies. This will present the greatest challenge, as examples of patterns and reuse in the AI-related technologies should be identified and favorably compared to the same in OO. Concerted efforts should be made to combine a "ring of familiarity" with a "nuance of innovation" in describing intelligent agent technology. Such a coupling of the old with the new will provide the cohesion that binds reality with possibility in the minds of the audience. This may provide a smoother transition into the use of the new technology if a level of comfort is provided with the initial introduction. The fact that Jeorg Mueller's book (The Design of Intelligent Agents) speaks of "patterns of behavior" ("PoB") and K. Gartner et al's book "Cognitive Patterns" (June 1998, Cambridge University Press) speaks of "cognitive patterns" already lends itself to the readily established OO model. The possibilities of extending Mueller's Forklift scenario to automated transportation systems imply a transfer of patterned behavior. Patterns of learning, adaptation, reasoning, perception, etc. could be addressed in an extension of UML specifically oriented towards development of agent-based applications. An agent does not necessarily correlate to a single object – it could be an application corresponding to several objects acting in unison to achieve a common goal. A one-to-one correspondence is not mandatory.

Establishment of a common forum of communication and understanding: conferences, academic mentoring, and metaphors

The chasm that currently exists between current and potential academicians, developers, and users regarding intelligent agents needs to be bridged. This could be accomplished through joint special interest groups, conferences, symposia, workshops, practicums, etc. Of utmost importance is seeking a common

denominator of communication and understanding. Since the intelligent agent concept itself appears to be a technical bridge between AAAI, OMG, FIPA, and their associated events (PAAM, Autonomous Agents, ATAL, AAAI, OOPSLA, ObjectWorld, Internet and E-Commerce, etc.), a common viewpoint may be relatively easy to achieve. This could result in a "cross-pollination" of sorts.

In the same vein, knowledge transfer can be achieved through cooperative relationships between academia and business. This could be accomplished through mentoring programs whereby in exchange for academic credit and monetary stipends, undergraduate and graduate programs could place students and Ph.D. personnel in "seedbed" companies that are just beginning their development efforts. Consulting companies that are profit-oriented may have a tendency to foster dependence instead of independence in the client, especially if it means future income potential. Such a scenario may be counterproductive for the success of intelligent agents, in that the client's options and exposure to alternatives may be constrained.

The transition of developers and users that come from a static and predictable computing environment can be facilitated through the use of metaphors. This marriage of the symbolic with the discrete could be accomplished by making frequent references to the area of human cognition: i.e., how and when these principles are applied in the area of intelligent agents. Comparisons and contrasts can be drawn between the discrete and symbolic computing paradigm, and one can be used to explain the other. For example, how ontologies can be viewed as an expansion of problem domains, how belief sets can use object states as an explanatory baseline, etc.

Encapsulation of complexity

The mainstream developer currently faces issues of managing budgets, political considerations, deadline constraints, a fluid and constantly changing environment, unpredictable user relations, etc. If a level of complexity is not readily justified in the development environment, then it will present a hindrance to a tool's productivity. In view of the complexity involved in neural nets, Bayesian networks, pattern recognition, and inference engines, it would behoove a vendor to hide these components inside a simple interface so to foster a more favorable learning curve. It is possible for a developer to understand the power and possibilities of a tool without comprehending its intricacies. If the community at large is initially alienated due to complexity, it's very hard to recover due to the lasting effects of first impression. This leads to the next strategy.

Maintaining an Identifiable "Input/ Ouput" Orientation

Currently, the power and capabilities of an object class or component is viewed in terms of services offered: input parameters, output results, and processing scenarios. This is in the spirit of "plug-and-play", whereby the insertion of a component in an application can be clearly identified and described in simple, functional terms. If the tool requires dispersion and embedding of its code in too many areas of the system, it may not be viewed favorably.

Conformance and tuning of the tool to the work environment of the targeted user or developer

Considerations of product use, features, scalability, vendor support, learning curve, training/onsite consulting, standards embrace, terms of purchase, etc. all need to be considered and addressed as part of the intelligent agent tool sets. These are traditional criteria of selection and acceptance, and play an important part in the "sales" job that both vendors and developers must perform in order to get the tool set purchased.

Five "P"s of Intelligent Agent Strategic Planning: Participation, Package, Pinpoint, Position, Publicize, and Penetrate
In the same manner that vendor products are vying for acceptance in the marketplace, technologies are also seeking the same kind of acceptance in the development mainstream. The many competitive offerings of technology (not necessarily products) may some day develop into a commodity-like marketplace, whereby acceptance of the offerings may depend more heavily upon their packaging and presentation than on performance or technical eloquence. For example can a UNIX daemon or a stand-alone expert system offer the same capabilities as an intelligent agent (N. Jennings and M. Wooldridge, Tutorial 5: "Applying Agent Technology" Agents '98 Conference)? The strategy that should be pursued in selling the concept of intelligent agents should be no different in nature than that of marketing a product to an identified audience while operating within a group of competitive vendors.

- *Participation:* the coherence of the strategic effort depends on participation and involvement of the academic, vendor, and targeted developer/user community. Agreement, cooperation, and synchronization of efforts will result in a properly orchestrated effort that will generate results. Fragmentation of the involved communities may be counterproductive to IA's initial acceptance.
- *Package:* define features, capabilities, advantages, relevant and practical application environments, etc. of IA.
- *Pinpoint:* define the optimal point of entry, ultimate desired area of emergence, targeted audience of users and developers.
- *Position:* differentiate IA versus alternatives (daemons, stand-alone processes), current advantages vis-a-vis competitive offerings, future advantages relating to changing environments, etc.
- *Publicize:* attendance at (and involvement in) the specific conferences, symposiums, workshops, publications, etc. that relate directly to the optimal point of entry and desired area of emergence (i.e., OOA/OOD/OOP, internet and cyberspace, network performance, etc.).
- *Penetrate:* Immediately following an initial "proof-of-concept" phase in the targeted area, addition AI-related tools can be brought in to enhance current product offerings. To the untrained developer or user, these capabilities may be perceived as rather daunting (and even traumatic) at first, but an effective communication strategy that utilizes metaphors and familiar terms and concepts could mitigate this adjustment. As the variation of tools and practical applications increases, penetration deepens accordingly.

Communicate initially in terms of the practical (not exclusively the potential) application of agents

The mention of phrases like "agent conflict", "virtual societies", etc. will not be interpreted in a familiar context, and may thus be construed as threatening. Framing realistic benefits of intelligent agents in terms that are familiar to the mainstream and can be viewed as benevolent will go a long way. As capabilities and offerings increase in number and become available, gradual indoctrination

can be applied. A successful proof-of-concept will automatically open the door for the subject of "potential" to be raised. The goal is convergence, not collision. This leads to the next strategy:

Gradual indoctrination of targeted population to the future possibilities of intelligent agent technology

This refers to the mental transition of a developer from a simplistic stimulus-response orientation in a predictable environment to an evolutionary approach in addressing application behavior. For example, instead of viewing an application as addressing current needs with periodic maintenance and enhancements, an IA solution can anticipate and plan a path of evolution whereby the application can adjust, learn, and autonomously resolve new environmental conditions without developer intervention. This is an unfamiliar territory for the mainstream developer, and will require careful training and "re-orientation" of the general approach application philosophy.

Capability of reducing the model and methodology to a "mental image" in the developer's mind

The lack of available resources and tight time frames cause a developer to maintain a "visual" image of procedure and practice in his or her mind. For example, with experience in the object-oriented arena comes the ability to mentally assemble and order objects, associations and relationships, paths of state change and their associated events, etc. without reference to strictly documented or formalized specifications. This is not the advised or even preferred method, but it still exists. Ideally, the agent model and methodology should be represented in relatively simplistic terms as a baseline, with detailed expansion offered as one's mastery progresses. It's this simplistic representation or mental model that may be referenced when time and resources are in short supply.

Vendors providing third party "plug-and-play" components or enhancements to current toolkits

The above issues of "Input/Ouput" Orientation and establishment of standards or guidelines will facilitate the use of component-oriented software: neural net, learning and adaptation, cognition, etc. The building of an enormous tool set that attempts to address and capture all needs of intelligent agent development may be too weighty and thinly spread for widespread acceptance.

Conclusion

There exists tremendous potential and opportunity in the area of intelligent agents if the proper obstacles are identified and addressed strategically by both the academic and vendor communities. The highest priority should be given to the establishment of an alliance with the target population, not its alienation. This could be performed through the pursuit of proper approaches and careful positioning of the technology. The software development field is in need of a revolution, and intelligent agents could very well be the first step. The emergence of any technology that may suffer from misunderstanding will require a unified effort of "push" (from the creator or academic community) and "pull" (from the receiver or vendor community).

Agent Technologies for Federated Information Systems in the Public Administration

Maurizio Panti, Aldo Dragoni, Luca Spalazzi, Claudia Diamantini and Salvatore Valenti
University of Ancona, Italy
{panti, dragon, diamanti, spalazzi, valenti}@inform.unian.it

INTERDATA is a research program, financed by the Italian "Ministero per l'Università e la Ricerca Scientifica e Tecnologica", devoted to the development of Internet-related technologies. One of the topics of this project is the automatic integration of information extracted from WWW sites. At the "Istituto di Informatica" of the "Università di Ancona", we are facing this problem in the context of interoperability of Information Systems (IS) for the Public Administration. This report deals with the IS for the "Piano di Inquadramento Territoriale" (Plan for the Integration of public and private Territorial entities) that will be an open and cooperative IS encompassing social, economical and cultural phenomena relevant for the local public administration.

IS*PIT is a complex, distributed and heterogeneous IS, determined by "political instances", hence designed upon pre-existing (and evolving) ISs. It cannot be conceived just as a network of interconnected computers, but it must include some mechanism of proactive information exchange and revision. IS*PIT is mainly characterized by the ability of integrating heterogeneous solutions adopted by different actors. It must be able to handle modifications due to both internal actions and external events. Features such as plurality, heterogeneity, dynamicity and reconfigurability, make traditional IS paradigms quite inadequate for IS*PIT, since the organization models they are based on (typically hierarchical or functional) are too rigid and too centralized in terms of the definition of languages, protocols, conceptualizations and control mechanisms.

Since the autonomy of each actor cannot be constrained to reduce heterogeneity, it is necessary to identify solutions that allow a reasonable level of information interchange without affecting autonomy. In the light of current research insights and available technology, an adequate model appears to be that of "federated information systems": that is a set of autonomous ISs which decide to share (part of) their data and processes in a cooperative perspective.

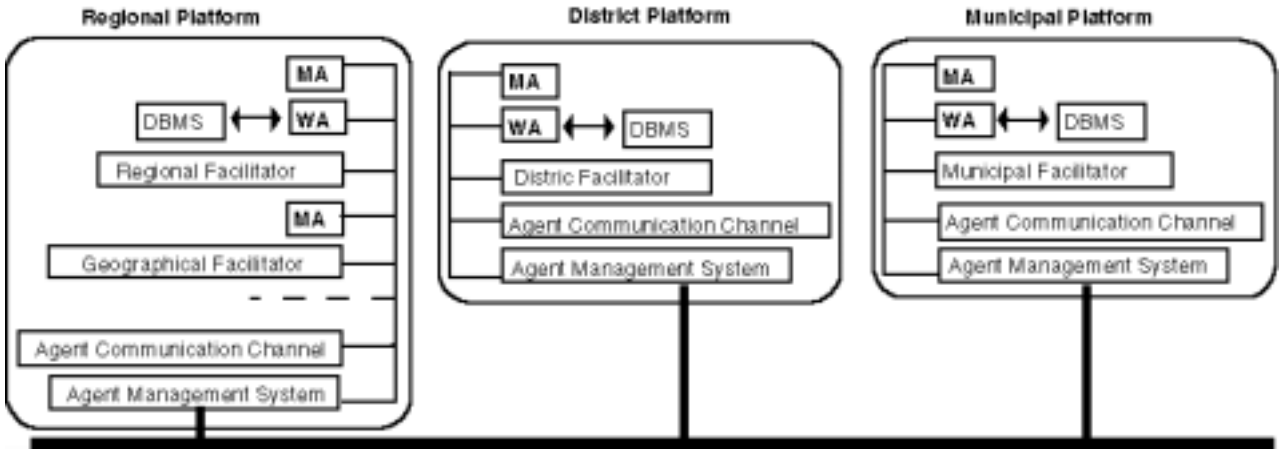

Figure 1. Architecture of IS*PIT: example of Platforms and Domains

The variety of actors concurring in PIT (public and private companies, associations, etc.) might be considered as provider/receiver of predefined information, as well as modifiers of information flows. They may propose plans, scenarios and goals, acting in different, but not necessarily disjoint, domains such as: political, legislative, economical, industrial, administrative, technological, cultural and formative. This charges the system with the burden of evaluating relevancy relative to each actor, while respecting constraints of quality, timeliness, accessibility and communicability.

Denoting "information repositories" the whole set of data stored and managed by each actor's IS, we distinguish the following categories of data/information:

1) structured and *available* in electronic format (e.g. files, records, and tables);

2) unstructured but available in electronic format (e.g. letters, resolutions, reports);

3) unstructured and *not* available in electronic format (e.g. hard-copies of documents).

Archives of category 1 are typically organized in files, directly managed by the operating system (sequentially, indexed, random, etc.), and only few recently systems take advantage of DBMS. Repositories of category 2 are usually spread on various computers and are often organized for individual purposes; their classification is usually semantically poor and "protocol-oriented". Finally, information of category 3 are not available at all, and their recovery is very limited and expensive.

Almost all actors have automated information systems and possess large repositories of data belonging to the first category, which could be shared in a cooperative perspective. Unfortunately, heterogeneity in data, information contents and organization structures hamper cooperation. These kinds of heterogeneity stem out from each actor's autonomy in structuring its own data repositories, processes and information flows. Unfortunately, in public administration systems, information repositories of categories 2 and 3 are very huge. Their automatic management is usually limited to record operations, in a pure storing perspective, making difficult to exploit them for decisional processes.

The Proposed Architecture

We stress the following features that IS*PIT should embody:

- *variety of information systems* from different "local entities"
- *actors autonomy* in the definition of its information needs, on what is to be regarded as public domain and in the design and implementation choices
- *dynamicity of the federation* which should permit new entries and/or the departure of old actors.
- *dinamicity of the information systems* which may change at the conceptual or at the physical level.

These points strongly affect the system design since:

- *it is impractical to build a unique integrated schema*
- *it is impossible to build global views and schemes "a priori"*.

These specifications led to *mediator-based systems* in the realm of agent-programming. Each agent is related to a *mediator* [1, 4, 5, 6, 7, 8] that is a *static*, *autonomous*, *intelligent* and *deliberative* agent. *Wrapper Agents (*WAs*)* translate queries in the local format, just as mediator agents do. WAs manage queries related to both data and schemas of the involved databases. Alike the wrappers defined by FIPA, our WAs are used to interface agent-based software with legacy systems. Each mediator is completely autonomous and might have different needs even belonging to the same organization. This justifies the need of a mediator for each agent of the system.

MAs build their conceptual structure by interacting with other agents and with human actors. They interact with WAs to access their original data schemas and with other MAs to access their views. A query submitted to a MA is redirected to entitled WAs if it refers to knowledge already owned by the MA. Otherwise, the MA has to enlarge its view in order to contain the knowledge required to forward the query. This may be obtained through the interaction with other WAs (to access their original data schemas) or other MAs (to access their global views). Thus the view of a MA is dynamically updated, strongly influenced by the queries submitted by the users. The interaction with other MAs can be useful both to create new views [2] and to improve the confidence on pre-existing views.
We decided to have an agent platform for each actor. This makes each actor free to manage his own agent and to select the services provided to other actors. A

number of different domains can be identified in the scope of IS*PIT. Actors co-operating in providing information in the domain of interest, are to be considered part of the domain itself. The same situation arises for the repositories hosting domain-specific information. This is the design rationale for adopting an agent domain for each sphere of interest. All WAs interfacing information systems involved in a domain and all MAs creating a view of the domain belong to the domain itself.

Agent-Prolog

As MAs and WAs should embody Artificial Intelligence technologies, we choose Logic Programming for its respectable tradition of useful tool for building inferential software agents. Besides Prolog, we obviously recognized the importance of Java as a platform-independent programming language. So, we realized the need of a Prolog interpreter written in Java as a first step towards a testbed laboratory for our software agents.

References

[1] Y. Arens, C. Y. Chee, C. N. Hsu, C. A. Knoblock, "Retrieving and Integrating Data from Multiple Information Sources", in International Journal of Intelligent and Cooperative Information Systems, vol. 2, pp. 127-158, 1993.

[2] C. Diamantini, M. Panti, L. Spalazzi and S. Valenti, Federated Information Systems Architecture for Local Public Administration, accepted for pubblication in in "Managing Information Technology Resources in Organizations in the Next Millennium", M. Khosrowpour ed., Idea Group, 1999.

[3] FIPA 97 Specification: Part 1-7, Foundation for Intelligent Physical Agents, October 1997.

[4] J. Hammer, J. McHugh, H. Garcia-Molina, "Semistructured Data: The TSIMMIS Experience", in Proc. I East-European Workshop on Advances in Database and Information Systems - ADBIS'97, St. Petersburg, Russia, Sept. 1997. Available at http://www-db.stanford.edu/pub/papers/adbis97.ps.

[5] T. Harder, G. Sauter, J. Thomas, "Design and Architecture of the FDBS Prototype INFINITY", in Proc. EFDBS97 - Engineering Federated Database Systems, Barcelona, June 1997.

[6] A. Levy, A. Rajaraman, J. Ordille, "Querying Heterogeneous Information Sources Using Source Description", in Proc. of the 22th Int. Conf. on Very Large Databases - VLDB'96, Bombay, India, 1996.

[7] V. S. Subrahmanian, S. Adali, A. Brink, R. Emery, J. J. Lu, A. Rajput, T. J. Rogers, R. Ross, C. Ward, "HERMES: A Heterogeneous Reasoning and Mediator System", submitted for publication, 1998.

[8] G. Wiederhold, "Mediators in the architectures of future information systems", in IEEE Computer, vol. 25, N. 3, pp. 38-49, 1992.

Just-in-time Information

Chris Knowles
CK Limited, Corporate Knowledge Software, UK
chris@cksoftware.co.uk

This article describes an application developed by CK Limited for the Corporate Finance department of a major Investment Bank, in which intelligent agent software is used to integrate information from external sources, and display the results on the user's desktop via the corporate intranet. The business information department within the bank is regularly asked to provide company profiles including annual report, latest news, company fundamentals, shareholder details and so on. All these items are obtained on-line from different sources. The CK1 Company Briefing Pack uses intelligent agent software to automate the process of searching across multiple sources of information and displays the results to the end user, on the desktop, in a way which is fully structured, managed, and accountable.

The recent white paper from the UK government's Department of Trade and Industry (DTI) on Competitiveness and the Knowledge Driven Economy, released on December 17th 1998 [1], gives the ultimate gloss of respectability within the UK to the field of Knowledge Management, which is now not only seen as crucial to the success of an individual company or corporation, but as a means of securing the economic future for the country as a whole.

The Financial Times, in its review of the White Paper [2] said that perhaps the most important aspect of the 'bundle of papers' was the economic analysis of the role of knowledge, set out in a separate 32 page analytical report. The Knowledge Economy is the new *big idea* which will guide policy in supporting British Industry through to the next millennium and beyond. Everyone may agree on the goal, but how to achieve it is still far from clear. As I wrote a year ago in June 1998 [3], the debate needs to shift from the "why" we need Knowledge Management, to the "how" to implement practical solutions that offer clear business benefit. This paper describes one such project, undertaken by CK Limited, in the financial services industry – one of the knowledge driven sectors identified in the DTI paper a spearhead for the future.

Just-in-time Information

Probably the biggest business development in the past 20 years has been the use of Just-in-time production, inventory, warehousing, delivery and logistics systems. This has resulted in greater efficiency, better customer service, better quality control, and massive savings, as organisations have reduced the financial costs of holding stocks of goods and raw materials, before they are required in the production process or for delivery to the customer. At the same time, customer service has improved, as organisations like Dell Computers manufacture in response to customer demand, build to order to meet a customer's individual specification, and still ensure delivery of

the finished product within 7-10 days, including shipping time.

In the Knowledge Driven Economy, information is the key commodity, and an increasingly important raw material for the production process. It is therefore somewhat ironic that, whilst organisations have been reducing their stocks of physical goods, they have been increasing their stocks of intangibles, and building more and more data warehouses, often at vast expense, to store the enormous quantities of information they are collecting as they go about their business. Many of the same principles which apply to physical materials apply equally to information. The right information, in the right place, at the right time, is just as important to the efficient operation of a knowledge intensive business, as the right component or raw material is to the efficient operation of a manufacturing production line. Information is expensive to acquire and difficult to store and organise. It gets easily lost or damaged in transit when passed from one person to another (as in the game of Chinese Whispers). Too much information is as bad, if not worse, than too little. Quality control is a nightmare, as information has a short shelf life and goes out of date as soon as you have stored it in your data warehouse, and just as with faulty raw materials, inaccurate information is not only useless but causes further damage.

Data storage is cheap, and data communications bandwidth is getting cheaper. As the DTI White Paper says, information is cheap and plentiful. Most commodities and raw materials like steel, oil or paper, are also cheap and plentiful. But if you are making a motor car, or stainless steel knives and forks, the right grade and quality of steel delivered when and where you need it, and quality checked in advance, costs rather more. This applies many times over to information. Delivered in bulk with no quality control it is worthless. Its value comes entirely from being relevant, accurate and up-to date.

Much information is unique to the organisation, and the efficiency with which it captures, stores, analyses, interprets, retrieves and displays the information is the mark of a successful business. But there is also a large and increasing quantity of information which is generally and publicly available. With the advent of Internet and Intranet technologies much of this is, potentially, available not just to the select few, but to everyone, on the PC or terminal on their desktop.

In a knowledge intensive industry such as financial services, organisations are prepared to spend large sums of money to ensure the rapid delivery of good quality, accurate, relevant and up-to-date information. There is a thriving industry of information providers, who make it their job to maintain the quality and integrity of the commodities and services they provide. An organisation such as Reuters invests large sums in ensuring the information they provide is up to date, accurate, and well structured and indexed.

Yet despite this, most companies still restrict access to information from external sources to the select few, or persist in a “heavyweight” approach, of taking data feeds, downloading and storing the information on their own internal systems. They do this despite the information going out of date as soon as they download it. In some cases they also lose the indexing and structure so carefully and expensively created by the Information Provider, and buy a search engine to recreate a second set of indexes themselves.

One reason for this is incompatibility between different systems, which makes it difficult for an organisation to conduct a single search across multiple different sources of information. The issue can be resolved, at some expense, with respect to internal systems, where everything is under the direct control of the organisation itself, but is an order of magnitude more difficult where external sources of information are involved.

The Internet and World Wide Web have provided a quantum leap forward in establishing a de facto standard for data communication (via the TCP/IP protocol) and for display of the results of searches (via the standard World Wide Web interface and the Netscape or Microsoft Internet Explorer browser on the desktop). But there still remain major incompatibilities at the application level in areas like password input, user access validation, search strategy and functionality, and indexing and structure of the content itself.

As an alternative to the “heavyweight” internal data warehouse, CK Limited adopted a more flexible approach using intelligent agent software to retrieve information directly from the source, when required, and automate the process of obtaining and displaying the right information when needed, in the right place, at the right time. Just-in-Time Information.

Intelligent Agents

Agents are one of the hottest areas in software development, and also one of the most overused and abused terms in the industry. Although it is generally agreed that agents share certain characteristics, there is no generally accepted technical definition of a software agent. In practice people can, and do call virtually any piece of software an agent. Agents do not form a well defined category in the same way as other types of software, for example databases, operating systems, or programming languages.

Software Agents are not new and have been around in some form for at least 5 years. Ovum, the market research company, in their report “Agents on the Web,” made the point that it is not the underlying technology that is new, but the way it is applied to the Web. What matters is not agents as such, but what they can do. Information Retrieval is the ideal field for Agent software. The problem of Information Overload is well known. How do you handle an enquiry on a web search engine that tells you have 535,798 hits? And however good the search engine is, at best it covers only one third of the ‘searchable web.’ [4]

Less well known, but equally common, is the problem of Intranet anarchy. This was graphically described by Udo Flohr in his article “Intelligent Intranets” in Byte magazine (August 1997) – ‘Intranets can be anarchy until you manage who can do what where.’ He describes two typical phases of Intranet development: ‘Step 1: Make a Mess’ followed by ‘Step 2: Clean it up.’

Every organisation is different, and has its own specific and individual requirements for information, based on its preferred way of conducting its own business. Two published sources of information may appear very similar, but an organisation will have very good reasons for preferring one to the other, based on the detailed content, coverage, structure and indexing, cost, commercial

terms, accuracy and relevance, ease of use, support and a host of other factors.

No one Information Provider can meet all the needs of all its potential customers. The result is Information Famine in an age of plenty. End users, with a PC on their desktop, may want to know something but often never find it, because it is too difficult, or takes too long, or they don't have permission to access the right database, or if they do, can't remember the ID and password.

The CK1 Company Briefing Pack uses agent software to automate the process of searching across multiple sources of information and displays the results to the end user, on the desktop, in a way which is fully structured, managed, accountable, and customised to meet the specific needs of both the individual customer and the organisation as a whole.

The CK1 Company Briefing Pack

In July 1998, CK Limited acquired all Intellectual Property Rights in the intelligent agent based web search products developed by Zuno Ltd. Zuno, a subsidiary of Mitsubishi Electric, was one of only four UK based companies, profiled by Ovum in 1997 in their report "Agents on the Web." The others were BT, Autonomy and ICL.

The CK1 Company Briefing Pack was developed initially to meet the requirements of the Corporate Finance department of a major investment bank. The business information department within the bank is regularly asked to provide company profiles including annual report, latest news, company fundamentals, shareholder details and so on. All these items are obtained on-line from different sources. The results of the various searches are printed, bound and distributed. An electronic version offers significant benefits over and above the printed version including:

- Automation of a largely routine process.
- The ability to search once across multiple sources
- Security – managed access to high value sources
- Cost control – all usage is logged and fully accountable
- Shareable – results can be accessed by anyone across the world
- Efficiency - cost and time savings
- Ease of use – results displayed in a standard format via the company Intranet

The greatest difficulty in automating this process arises because, when considered in detail, every information source is different, and the same search strategy can not be applied in every case. The problem is compounded because not only are the Information Sources different, but different search strategies are appropriate for different types of companies. For example the best sources for information for a US company may be different from the best sources, for the same type of information, for a UK or European company. In other cases 'smart search' strategies are required which, for example check two sources and retrieve data only from that which is most up to date, or modify the search based on an initial enquiry.

The CK Limited solution involves designing a site agent for each source, to handle the specific aspects of connection, input of ID and password, submitting the search and retrieving results. In addition search agents handle the smart search strategies, and a display or user interface agent manages the display of results to the user.
In overview, the application comprise three elements:

a) A display and access module (or agent), developed on a bespoke basis for each customer.

b) The CK1 "engine" which performs the generic searching, profiling, matching and filtering functions involved in making an enquiry and consolidating and displaying the results, including smart search agents.

c) Resource adaptors or 'site agents' for each supported information source, which handle aspects such as connection, security, password input and validation, navigation, retrieval and filtering of the appropriate result(s) for each source.

Suitable agents to support most of the major Business Information sources including Reuters Business Briefing, (News), Dow Jones Interactive (News), Disclosure Global Access, (Annual reports), FT Referencer, (Annual reports),

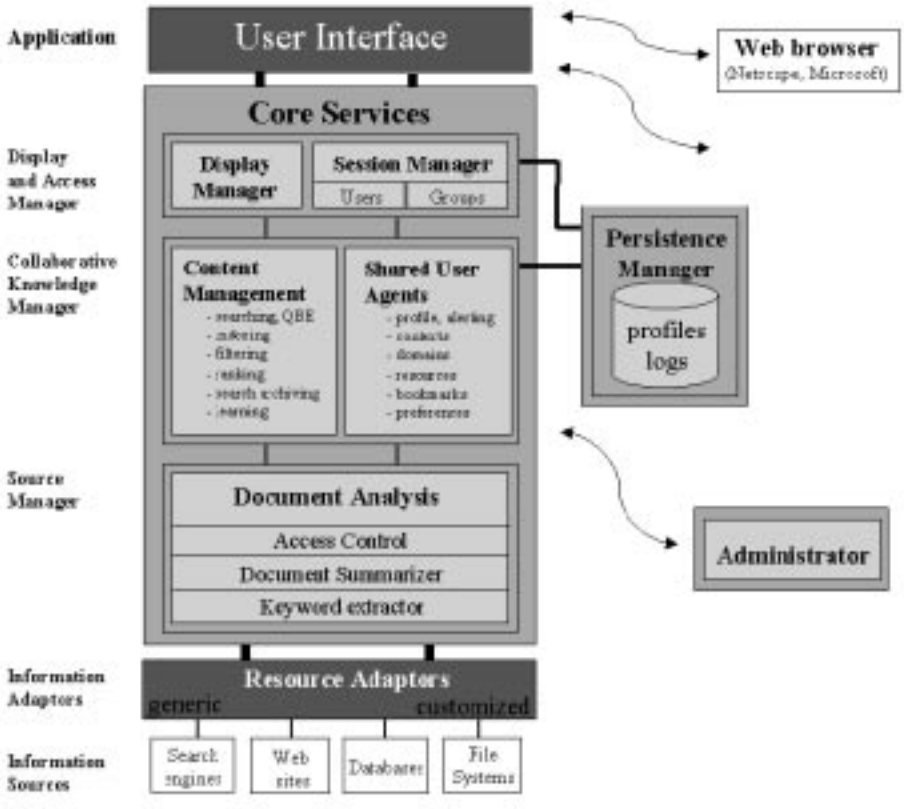

Figure 1: System architecture

Extel Cards, Investtext Research Bank, (Brokers' Notes), Citywatch, (Shareholder information), and TED (The Estimate Directory) have been developed, or are currently in development. The system is written in Java, using Java servlets, on a JavaWebserver platform running under Microsoft NT.

During development, a number of issues, both technical and commercial, arose and were either resolved or will be addressed in future releases. Some of the more significant of these were:

- The Commercial Relationship with the Information Provider.
- Diversity of company classification schemes.
- Smart search strategies
- The need for printed output
- Multi-site operation

Conclusion

Many organisations build databases of information, which can be accessed in a structured way via the corporate intranet. But this can push back the problem of maintenance and updating the information one level. Someone still needs to update the database. For internal information, a company has to take on the responsibility itself for building and maintaining its own store of corporate knowledge, using Lotus Notes, an internal document management package, or one of the many other intelligent intranet database management tools now available. But in the case of external information, why recreate what others have already done? The information manager for a major pharmaceutical company said at a recent conference in response to a question on this issue: "We are a drug company, not an on-line host. Wherever we can we buy. We don't want to create our own large databases of publicly available information."

An intelligent agent based "Just-in-time Information" solution, such as the CK1 Company Briefing Pack, provides the best of both worlds. It allows the company to select, manage and control the quality of the information displayed on the end-users' desk top and provides a highly personal and customisable service for the end user, but leaves all the problems of maintaining, structuring and updating the information itself to those who do it best, the Publishers and On-Line Service Providers.

References

[1] White Paper - Our Competitive Future, Building the Knowledge Driven Economy The Stationary Office, December 1998, or www.dti.gov.uk/comp/competitive

[2] 'Knowledge is Power' emerges as big idea for the economy, Financial Times, December 17th 1998.

[3] Knowles, Chris and Ferguson, Innes, 6 Techniques for better matching, filtering and profiling stored knowledge, Knowledge Management, Issue 6 vol 1, June/July 1998

[4] Lawrence and Lee Giles, Searching the World Wide Web, Science Magazine, p100, April 1998. Estimated figures for the fraction of the 'indexable web' covered by the major search engines are given as Hotbot 34%, AltaVista 28%, Northern Light 20%, Excite 14%, Infoseek 10% and Lycos 3%. The 'indexable web' excludes subscription based and password protected services.

Announcement

Turning the Global Infrastructure into a "Universal Information Ecosystem"

A strategic research initiative of the new European programme on "Future and Emerging Technologies"

Over the years to come, an ever-growing segment of the world population will be connected to the global information infrastructure and many aspects of life will increasingly take place through or within it. A complex, dynamic and heterogeneous infrastructure environment will emerge, where "infohabitants" (people, organisations, as well as virtual entities acting on their behalf, smart appliances, smart buildings, etc.) will communicate, cooperate, negotiate, self-organise, create and offer new intermediation services, taking advantage of a dynamic set of information and communication resources to achieve their objectives or those of their owners.

"Universal Information Ecosystems" (UIE) is a research initiative by the "Future and Emerging Technologies" part of the European programme in "Information Society Technologies" (http://www.cordis.lu/ist/). It stems from the vision of an open global infrastructure environment that constantly evolves in an attempt to best meet the changing demands of its dynamic population of "infohabitants". This would be an environment where, at any point in time, the knowledge and capabilities of every single infohabitant could be effectively enhanced and dynamically recombined with that of all others in an efficient manner, and where new infohabitants with new demands or offering new services could be seamlessly introduced. For this highly decentralised and heterogeneous environment to be as responsive as possible, its infohabitants would need at any point in time to be globally, yet selectively, aware of (or "sense") the opportunities in the environment that are most relevant to their objectives and activities. Adaptation to changing condi-

tions would take place through the decentralised creation, deletion, reconfiguration, migration and/or recombination/reorganisation of infohabitants and/or through the activation of self-regulating mechanisms in the ecosystem.

The process of reacting and adapting to a constantly changing environment will always in part be triggered by decisions directly made by people or organisations. In order to take full advantage of the emerging information ecosystem, UIE seeks to explore ways of automating parts of this process and to effectively support the delegation and realisation of many of these decisions.

Significant progress towards this vision of a Universal Information Ecosystem demands going beyond incremental extrapolations of current technological paradigms. UIE will seek to support research projects aimed at exploring and validating new technologies and scenarios for a Universal Information Ecosystem. Success is expected to depend on taking a broad and interdisciplinary perspective, pulling together expertise from areas as diverse as for example, distributed systems, software engineering, computational logic, artificial intelligence and human computer interaction as well as economics, life sciences, organisational theory or fundamental social science. By breaking across traditional boundaries, the initiative also hopes to help foster the creation of a new trend-setting research community to lead Europe to the forefront of this interdisciplinary area.
For more information on the initiative, please contact either:

Leonardo Flores
Tel: +32(2)299-1825
Fax: +32(2)296-8397
Leonardo.Flores @cec.be

or

Loretta Anania
Tel: +32(2)296-3491
Fax: +32(2)295-6937
Lan@postman.dg13.cec.be

Fourth International Conference on Autonomous Agents (Agents 2000)

Barcelona, Spain, June 3-7 2000

Autonomous agents are computer systems that are capable of independent action in open, unpredictable environments. Agents are also one of the most important and exciting areas of research and development in computer science today. This technology is currently being applied in domains as diverse as computer games and interactive cinema, information retrieval and filtering, user interface design, electronic commerce, and industrial process control. The aim of the Agents 2000 conference is to bring together researchers and developers from industry and academia in order to report on the latest scientific and technical advances, discuss and debate the major issues, and showcase the latest systems. The conference will include technical presentations of papers and videos, panel sessions involving internationally recognised experts in the field, software and robotic agent demonstrations, and an exhibits session. The conference will also include workshops and tutorials.

Conference Deadlines

October 25, 1999
Paper and video submission deadline.

December 6, 1999
Workshop and tutorial proposals due.

March 3, 2000
Software and robotics demo submission deadline.

March 17, 2000
Camera-ready copies of accepted papers due. Tutorial material and workshop papers due.

June 3-4, 2000
Workshops and tutorials.

June 5-7, 2000
Conference technical sessions.

General Chair:

Carles Sierra (IIIA-CSIC, Spain)

Program Co-Chairs:

Maria Gini (University of Minnesota, USA), Jeff Rosenschein (Hebrew University, Israel)

Senior Programme Committee

Jeffrey M. Bradshaw
(The Boeing Company, USA)

Marco Dorigo
(Université Libre de Bruxelles, Belgium)

Ed Durfee
(University of Michigan, USA)

Nick Jennings
(Queen Mary & Westfield College, UK)

Lewis Johnson
(University of Southern California, USA)

Hiroaki Kitano
(Sony, Japan)

Ramon López de Màntaras
(IIIA-CSIC, Spain)

Jörg Müller
(Siemens, Germany)

Pablo Noriega
(INEGI, Mexico)

Steve Smith
(Carnegie Mellon University, USA)

Katia Sycara
(Carnegie Mellon University, USA)

Manuela Veloso
(Carnegie Mellon University, USA)

Mike Wooldridge
(Queen Mary & Westfield College, UK)

For more information, vist:
http://www.iiia.csic.es/agents2000

Autonomous Agents

Fourth International Conference on MultiAgent Systems (ICMAS 2000)

Boston, USA, July 7-12 2000

Multiagent systems are computational systems in which several agents interact. The concept of "agent" as a computational entity that can accomplish goals has spawned a plethora of activity in artificial intelligence and more widely in computer science. ICMAS is the premiere international forum specifically devoted to the study of agent interaction and the development of multiagent systems. ICMAS participants typically come from the physical, computational, and social sciences, and are drawn together by their mutual interest in understanding the phenomena that arise when agents interact, in investigating the interplay between agents as individuals and as participants in collective settings, and in formulating languages, architectures, mechanisms, and engineering and evaluation methodologies that apply to multiagent systems.

ICMAS 2000 TOPICS include, but are not limited to:

Interaction capabilities, constraints, and preferences :

- Communication languages and protocols
- Organization and social structure
- Teamwork and cooperation
- Decentralized systems
- Mechanism design

Reasoning about coordinated interactions:

- Conflict resolution and negotiation
- Multiagent planning
- Coalition formation and organization self-design
- Agent modeling and plan recognition
- Multiagent learning
- Distributed search and constraint satisfaction
- Foundations (multiagent logics, game-theory, economics, philosophy, etc.)

Engineering, deploying, and evaluating multiagent systems:

- Agent programming languages
- Multiagent programming frameworks
- Agent models and architectures
- Standards for multiagent technology (interaction protocols, languages)
- Development and engineering methodologies
- Evaluation of multiagent systems
- Testbeds and development environments
- RoboCup, multiagent robotics, RoboCup-Rescue
- User interfaces and personalizable agents

Practical applications:

- Agents in electronic commerce
- Cooperative information systems
- Distributed resource allocation
- Information agents on the internet
- Multiagent simulations of social and biological systems
- Multiagent vision and robotics
- Believable agents in multiagent settings
- Interacting personal digital assistants

Important Dates

Nov 15, 1999 Deadline for electronic title pages & abstracts

Nov 17, 1999 Deadline for final papers Jan 10, 2000 Author notifications will be mailed

July 7-9, 2000 Federated Workshops & Tutorials at ICMAS 2000

July 10-12, 2000 ICMAS 2000 main conference.

Federated Events

ICMAS'2000 will feature special sessions and events highlighting some of the most exciting multiagent research issues and applications. These are expected to include multiagent robotics (with emphasis on RoboCup and RoboCup-Rescue) and trading in electronic markets. Other federated events and workshops (such as the Agents, Theories, Architectures and Languages ATAL 2000) are also planned at times adjacent to ICMAS.

Dates and Venue

ICMAS 2000 and its federated events will be held July 7-12, 2000. July 7th-9th are reserved for federated workshops and tutorials, while the main conference will be held July 10-12, 2000. The conference venue is the Boston Park Plaza Hotel, Boston, Massachusetts, USA. For further information, please see

http://www.bpph.com/.

Organizing Committee

General Chair:

Ed Durfee

durfee@umich.edu

Program Co-chairs:

Europe/Africa: Sarit Kraus

sarit@cs.biu.ac.il

Americas: Milind Tambe

tambe@isi.edu

Asia/Pacific Rim: Hideyuki Nakashima

nakashim@etl.go.jp

Publicity Chair: Christian Lemaître

cll@xalapa.lania.mx

Publication Chair: Piotr Gmytrasiewicz

piotr@cse.uta.edu

Finance Chair: Mike Huhns

huhns@sc.edu

Workshops Chair: Klaus Fischer

Klaus.Fischer@dfki.de

Tutorials Chair: Munindar Singh

mpsingh@eos.ncsu.edu

Regional Liaison for Asia/Pacific: Chengqi Zhang

chengqi@deakin.edu.au

Regional Liaison for Europe/Africa: Mike Wooldridge

m.j.wooldridge@qmw.ac.uk

For more information, vist:

http://icmas.lania.mx

What's Happening in AgentLink?

Mike Wooldridge
AgentLink, Queen Mary and Westfield College, UK
M.J.Wooldridge@elec.qmw.ac.uk

Feedback from April SIG Meetings

The April SIG meetings, held in London in concert with the PAAM-99 conference, were attended by about 160 people in total, of which only about 50 asked for AgentLink funding. The fact that AgentLink was only required to fund one in three participants is an impressive indicator of how valuable these meetings are perceived to be. As well as funding participants to attend, about 50 people took advantage of AgentLink funding to obtain free registration at the PAAM conference itself.

Next SIG Meetings

The dates for the next SIG meetings have been agreed as follows - either check out the AgentLink WWW page for latest information, or contact the coordinator directly:

- **Coordination and Control**
  Barcelona, Spain,
  6-7 July 1999
  Coordinator: John Perram
  (jperram@mip.ou.dk)

- **Intelligent Agents for Telecoms Applications**
  Stockholm, Sweden,
  11 August 1999
  Coordinator: Sahin Albayrak
  (sahin@dailab.cs.tu-berlin.de)

- **Agents and Electronic Commerce**
  Barcelona, Spain,
  20-23 September 1999
  Coordinator: Carles Sierra
  (sierra@iiia.csic.es)

- **Software Engineering and Methodologies**
  Barcelona, Spain,
  20-23 September 1999
  Coordinator: Jan Treur
  (treur@cs.vu.nl)

- **Intelligent Information Agents**
  Barcelona, Spain,
  20-23 September 1999
  Coordinator: Matthias Klusch
  (klusch@dfki.de)

- **Agents for Simulation**
  Barcelona, Spain,
  20-23 September 1999
  Coordinator: Rosaria Conte
  (rosaria@pscs2.irmkant.rm.cnr.it)

The Summer School

Much to our suprise, all 160 places on the summer school were filled by mid-May - two months before the summer school itself, and six weeks befoe early registration closed! AgentLink is supporting about 50 students from AgentLink member nodes to travel to the event, and the remainder are self-funding. The summer school is offering 17 courses in total, and promises to be one of the summer's main agent events.

Free Subscriptions to Agents Jounal

AgentLink is purchasing subscriptions to the International Journal of Autonomous Agents and Multi-Agent Systems, published by Kluwer Academix Press, for all AgentLink member nodes. These will be sent to the technical contact point at each AgentLink node - expect your first copy soon! Our intention is to continue subscriptions throughout the lifetime of AgentLink.

Joint AgentLink/i3net Meeting

The long-term technological goal of the ESPRIT Long-Term Research pro-active initiative in intelligent information interfaces, i3, is to build novel human-centred interfaces for interacting with information, aimed at the broad population. The i3 initiative involves a supporting network of excellence, i3net, which includes a special Interest Group on Agents. While this SIG appears to be dormant, judging by the i3-conference in Denmark, there is considerable interest in agent technology amongst the i3 community. To enable i3net members and AgentLink members sharing common interests to meet, AgentLink and i3net together sponsored a workshop in Valencia at the end of June. The workshop, organised byJeremy Pitt from Imperial College, UK, was attended by about 30 people, and may be followed by further events if there is a demand.

AgentLink Travel Support

AgentLink made travel support available for authors with accepted papers at the CIA-99 workshop (Uppsala, July/August 1999) and the AMEC workshop (co-located with IJCAI-99 in Stockholm, August 1999). Contact coordinator@agentlink.org for details.

Agent Mediated Electronic Commerce - Second Meeting

Carles Sierra
Artificial Intelligence Research Institute, CSIC, Bellaterra, Spain
sierra@iiia.csic.es

After the kick-off meeting in Brussels last year this meeting was aimed at focussing on those particular topics that appeared to be of maximum interest to the research and development groups of the SIG. These topics both structured the meeting and proved to impact on the practical outcome of the meeting as explained below.

The meeting's overall objective was to discuss an initial assessment of a technological roadmap for the area of interest. This technological roadmap will become, as a result of the discussions of this meeting and the necessary post editing by all participants, the practical outcome of the meeting. This paper will cover one of the initial objectives of the SIG: to come out with a 'green paper' of the area. Hence, all participants had in mind the three key questions to answer during the meeting: where is Agent mediated Electronic Commerce going to be in the near future?, Where is agent mediated electronic commerce now?, and how are we going to make the way through?.

We organised the meeting around four working groups that attempted both to put on the table techniques that the European groups are working on and to answer the questions that would structure the technological roadmap. This report describes, in a cursory fashion, the discussions that took place and the topics that were addressed. The detailed paper, containing the concluding results, will be made available through the web page of the SIG by the end of summer.

The four topics found most relevant during the kick-off meeting were: negotiation, market and protocol specification, preference modelling and security. Before the meeting all participants showed their interest in one or several topics. As preference modelling and security showed a lower interest but had some arguable relation, we ended up with three different working groups. During the last afternoon a joint panel between the AmEC and the I2A SIGs took place.

Negotiation

The working group on negotiation started by an introductory lecture given by Nick Jennings. He argued that the second generation of e-commerce systems will be based on a much richer interaction than the current systems. This interaction would be based on argumentation and flexible protocols. A series of short presentations by Eugenio Oliveira, Carles Sierra, Nir Vulkan and Shamina Paurobally focussed next on different relevant aspects of the research on negotiation. The need for adaptive agents, special architectures for argumentative negotiation, and formal models for negotiation were the main topics addressed by these talks. Basically the agreement was reached that we need much more flexible forms of negotiation to overcome the currently very simple forms (such as auctions). Also, in order to be able to do negotiation on goods of high value, more efforts must be spent on better forms of trust modelling.

Finally, a panel chaired by Frank Dignum and with the participation of Nick Jennings, Ana Paula Rocha and Nir Vulkan discussed about the characterisation of negotiation problems, about which domains were more adapted for the application of negotiation techniques and about the series of problems to overcome. The discussions were very lively among the participants. It was clear that negotiation is probably the highest topic of interest to the SIG participants.

Market and Protocol Specification

This working group started with a detailed presentation on 'Markets and protocols' given by Frank Dignum. He covered many different aspects such as the impact of electronic markets on the supply chain or the role of organisational perspectives (institutions) in the design of markets. He extensively discussed the different types of auctions and argued that one of the hot topics of future research will be the design of electronic institutions. Then some short presentations by Maria João Viamonte, Fredrik Ygge and Carles Sierra focussed on mechanism design, electronic institutions specifications and on models for electronic market places.

Finally, another panel, chaired this time by Chris Preist, focussed on the possible alternative designs for markets. (Elements like reputation agencies and strongly institutionalised systems enforcing behaviour appeared as near yardsticks in the design of markets.) Travel agencies and retailing were identified as the domains where the next wave of applications will appear. As important factors to take into account when designing markets, efficiency and trust were highlighted. It was suggested that a good application area to focus was that of high liquidity markets. The panelists expressed their opinion that in order to design market more research is needed at the frontier with marketing, psychology and social simulation.

Preference modelling and security

Vania Conan talked about the problem of preference modelling within the context of the European funded project AIMEDIA. He argued about the importance that personalization has in the context of electronic commerce and proposed that machine learning techniques are needed for personalization to be more effective. In particular, data mining techniques seem very appropriate to determine users' profiles. The other two big research areas identified by him were negotiation and context perception. Mehdi Dastani and Gerd Voelksen gave short presentations about the important issues in preference modelling, i.e. representation, learning and roles.

A panel, chaired by Markus Schwehm, focused on different questions about preference and security. The participants in the panel were Gerd Voelksen Mehdi Dastani and Vania Conan. The business scenarios detected as the most promising for the near future were business-to-business and retailing. Panelists argued that preferences should be modelled as a dynamic component, hence the need for learning and case-based reasoning techniques. Also, fuzzy logic and data mining where considered as interesting general approaches for profiling. The open research issues mentioned included standards for security and the automatisation of user profiling.

Joint panel AmEC-I2A

In the afternoon of the second day a joint panel between the I2A and AmEC SIGs took place. Walter van de Velde was the chair and the participants were Markus Schwehm, Eugenio da Costa, Vania Conan, Fredrik Ygge, Mihai Matskin and Matthias Klusch. The discussion was centred on the intersection that can be found between the integration of physical and immaterial spaces, the technology, and the user. Many of the arguments that appeared during the previous discussions reappeared here, namely negotiation and adaptability. A paper containing the detailed discussions will be edited by Walter van de Velde. Check the web pages of the SIG to keep track of the status of the paper.

Analysis

The atmosphere of the discussions in this meeting was very good. I think that we made quite a progress in detecting which are the relevant points for the research in the upcoming period. Thus, my assessment is that we are in a position where a clear roadmap can be prepared among the participants. I'll take the responsibility of editing the paper from the input received from the participants in the meeting.

We also agreed that the next meeting will take place in Barcelona during the week 20-23 of September. I'll take care of the arrangements for this meeting. The focus will be around the final discussions on the draft of the roadmap and the presentation of papers to be included in the book that we will prepare as a general outcome after the first year of work of this SIG. The book will be organised around the main topics detected in the roadmap.

More information on this SIG can be found at:

http:///www.iiia.csic.es/AMEC

Methodologies and Software Engineering for Agent Systems - Second Meeting

Jan Treur
Department of Artificial Intelligence, Vrije Universiteit Amsterdam, The Netherlands
treur@cs.vu.nl

This SIG meeting was organised in the form of four half-a-day working groups. Themes addressed in the different Working Groups were:

- Design and Specification; Unified Agent Modelling Language
- Logical Foundations
- Requirements Engineering and Verification
- Relation to Existing SE Methodologies

Although the working groups in principle were independent, the participants in the working groups had a large overlap, with an exception of the logical foundations working group, which had a substantially lower number of participants.

In each of the working groups a research agenda for the theme was discussed (in continuation of the first SIG meeting), for the short term (1999), mid term (2000 - 2001), and long term (> 2001).

Design and Specification and UAML

The chairman of this Working Group was Jean-Pierre Müller. Presentations were performed by: Marie-Pierre Gervais, Bernhard Bauer, Rune Gustavsson, Olivier Gutknecht and Jean-Pierre Müller. The discussion focused mainly on the development of a Unified Agent Modelling Language (UAML).

The following short term aims were discussed:

- Define requirements for UAML (e.g., precise semantics, verification and validation, automatic code generation, include primitive agent concepts; which
- Identify what is already achieved by UML, and what are shortcomings (e.g., complex behaviour requirements, knowledge-intensive functionality)
- Define a first extension of UML
- Test this extension in projects
- Submit a proposal to FIPA2000

Identified mid term aims are:

- Develop graphical language and modelling tools
- Develop formal semantics
- Develop code generation tools
- Develop execution, test, debug tools
- Integrate both emergentist design and agent-based design methods
- Develop precise relations between global (organisational) specifications, local external agent specifications, and local internal agent specifications
- Identify and characterise classes of applications (e.g., distributed manufacturing systems, information agents, Electronic Commerce, ...)
- Develop dedicated modelling elements for a given class of applications
- Develop libraries of reusable models/ patterns for: organisations, interactions, reusable agent models, agent compo-

nents, tasks/problem solving methods, ontologies, knowledge bases

- For a given implementation platform, develop guidelines/automated support for implementation of a model

Long Term aims:

- Further maintenance and development of the modelling language
- Supporting tools (A-CASE tools)
- Development of formal semantics

Logical Foundations

The Working Group on Logical Foundations was chaired by Michael Fisher. Presentations were performed by him and John-Jules Meyer. It was felt that the use of logical foundations in practice is problematic, and it is not clear how this can be changed. A number of difficult questions were identified:

- What use are logical foundations in practice ?
  (of no use in today's practice, should become of more use (verification), but how ?)
- How can we get more people to use formal approaches ?
  (visualisation & graphical tools ?)
- Will a logical semantics of UAML be possible ?
- Is a grand unified (logical) theory possible ?
  (is far away)

Requirements Engineering & Verification

Chairman of the The Working Group on Requirements Engineering and Verification was Jan Treur. Presentations were performed by Michael Petit, Jan Treur, and Paul Kearney.

As a short term aims it was identified:

- Development of a software environment for verification based on theorem proving

Mid term aims identified were:

- Development of a software environment for verification based on model checking
- Informal, semi-formal and formal representations of requirements and scenarios and relations between them (graphical form, library of requirement templates/patterns)

Long Term aims identified are:

- Two way relation between global phenomena and properties of agents
- Development of a software environment for requirements engineering for agent systems

Relation to Existing SE Methodologies

The Working Group on Existing SE Methodologies was chaired by Carlos Iglesias. Presentations were performed by him and Alexander Lauert.
Short term aims identified are:

- Determine demands specific for agent development methodology
- Compare these demands to offers from existing SE approaches

Mid and long term aims:

- Better exchange and integration with other communities (e.g., SE/OO, KE, RE, PE, DCS, logic-based)

General issues

Some general issues put forward are not restricted to one of the Working Groups, but of a more general type:

- Determine the different models that are needed (e.g., organisational specification, requirements at different process abstraction levels, system design specification)
- Relations between different forms of specification
- Create glossary of agent concepts
- Characterise what makes agent systems different from other systems
- Relations between formal and informal
- Reusable ontologies (in XML ?)

On the basis of the identified issues a road map for the area of this SIG will be designed this year. A draft version of this road map will be discussed at the third SIG meeting in September in Barcelona.

More information on this SIG can be found at:

http://www.cs.vu.nl/~treur/SIG.meth0.html

Agent-based Social Simulation

Rosaria Conte
Division of AI, Cognitive and Interaction Modelling, IP/CNR, Rome, Italy
rosaria@pscs2.irmkant.rm.cnr.it

The social simulation research community has developed rapidly in recent years, attracting sociologists, computer scientists, logicians and economists (and this list is not exhaustive) who believe that simulation methods can support the analysis of social issues while maintaining the relevance of sociology and the rigour of economics. The result is a happy research community in which participants from different discipline backgrounds use each others' techniques and approaches to develop new methods and methodologies. We have seen the use of formal logics to represent such concerns of sociologists as trust, belief and helpfulness. Results and representations from cognitive science are regularly used to specify agents in computational organisation theory and a new approach to economic issues unfettered by old, conventional approaches.

This SIG is intended to reinforce the development of a strong European presence in social simulation and also to demonstrate that the disciplines on which the social simulation community draws can themselves benefit from the results obtained by members of that community. Perhaps because of its relative youth or because it encompasses an exceptionally diverse set of academic backgrounds, the social simulation community remains

open to new ideas and to new applications. This is a strength on which the SIG is intended to build. The SIG is not intended to be another home for conventional economics. Applications of equilibrium representations of markets and game theoretic approaches to, for example, electronic commerce are already widespread. Where information markets and electronic commerce are applications of social simulation, a bottom-up modelling approach will be encouraged. There is already at least one such project arising from the activities of AgentLink and others are to be encouraged.

The overall aim of the activities of the Agent-Based Social Simulation (ABSS) SIG is to support and initiate joint projects and collaborations among active SIG participants towards research on and development of social simulation for the support of social analysis and to inform developments in software engineering. The SIG supports these aims by

- putting groups with related interests in touch with one another
- providing a discussion forum via regular SIG meetings.
- disseminating information about developments in the field of social simulation
- organising meetings around existing events with topics strongly related to agents and social simulation, such as interdisciplinary workshops, conferences, and symposia.

The First Meeting

The first meeting was characterised by a multi- and inter-disciplinary audience (over 40 participants), including among others Multi-Agent Systems and cognitive scientists, environmental scientists, logical philosophers and philosophers of law, sociologists, economists, physicists, operational research scientists, psychologists.

The meeting was inspired to the principle of promoting a general discussion around research projects and innovative ideas. Presentations and discussions were divided in paper sessions and one panel discussion session. Short presentations have been selected to stimulate discussion in the following main domains:

- *Agent Technology*: existing platforms and architectures in the field of agent systems, which are or can be used for agent-based(social) simulation platform. In this session, some projects for addressing issues of both social scientific and MAS interest were presented and the possibility to implement simulation models with agent architectures rather more complex than those usually employed in the social simulation community were discussed (for example, agents endowed with beliefs, desires, and intentions).
- *Agent Modelling*: the level of complexity of the agent model has been addressed. Some speakers argued that autonomous intelligent agents ought to be implemented in order for simulation models to deal with such social and institutional phenomena as the *violation* of norms, the *solution* of normative conflicts, the *acknowledgement* of norms and institutions, the *formation* of new ones, etc..
- *Modelling social institutions and group formation*: in some contributions, a radically simplified notion of agency was applied to the study of the emergence of institutions within the framework of complex dynamic systems. On the other side, contributions from the logical philosophical field showed the necessity to integrate a formal/theoretical approach in producing a consistent and heuristic institutional ontology -what is a norm, how to distinguish and therefore acknowledge rights and permissions, etc.? These questions were found preliminary to the modelling of social systems. Finally, the role of social emotions in group performance was also addressed.
- *Infrastructures and Tools*: several different systems and languages have been presented (DESIRE, SDML, MASSIF, CORMAS, etc.). Although not necessarily a disadvantage, this state-of-the-art dissemination requires a common investigation of similarities and specificities of existing tools.
- *Applications*: this range from agent-based simulation of MAS for purposes of evaluation, to participatory policy-making both at the level of development and at the level of verification of the candidate policies. In areas such as resource management, where physical and socio-economic factors are strictly intertwined (Integrated Assessment Modelling), the agent-based approach to social simulation was found especially useful.

During the panel discussion session, the impact of the field on (social) theory building, testing, and its application was emphasised. The issue of how to validate simulation findings, and the necessity for confrontation with adjacent fields (e.g., cultural evolution and neural nets) were identified as possible foci for future meetings.

One fundamental feature of social phenomena was generally acknowledged to consist of feedback loops (for example, between agents to institutions and from institutions back to agents). At the application level, this assumption has a direct impact on the role attributed to agent modelling and implementation in the field of participatory policy-making. Another largely shared assumption concerns the insufficiency of the model of economic rationality as a model of agency for social simulation.

Future work

A number of open questions were identified:

- At the theoretical-conceptual level, the issue of emergence was re-examined since its definitory and operational criteria are not clear yet (the relation between emergence and evolution, the direction of emergence - whether from lower to higher level of description or both-, etc.).
- At the epistemological level, two different principles have been debated at some length: the so called “KISS” principle (Keep the model Simple, Stupid) Vs the “KICS” principle (put Conflicts and concurrence into it) Vs the “KIN” principle (model Natural phenomena). The necessity to refer to real and natural phenomena Vs artificial but plausible ones was also discussed.
- At the methodological level, the question of validation is an issue of major concern for social simulators. What kind of data are simulation data?
- At the application level, the utility to design modular tools, which can be implemented one on top of another, Vs concurrent ones and the respective advantages of either option was also discussed.

More information on this SIG can be found at:

http://www.cpm.mmu.ac.uk/abss/

Intelligent Information Agents (I2A) - Second Meeting

Matthias Klusch
DFKI GmbH, Deduction and Multiagent Systems Lab, Saarbruecken, Germany
klusch@dfki.de

The purpose of the second meeting was threefold: Firstly, to report on the proceedings of collaborations which have been initiated and set up among SIG participants since the first meeting, secondly, to present new advanced work of R&D groups active in the I2A area, and finally, to discuss potential working groups within the SIG focussing on few main topics and significantly relevant issues of research and development of information agents. Accordingly, the agenda of the two-day meeting comprised two sessions for selected presentations of projects and groups, one session for software demonstration, a keynote talk, one common session with the AgentLink SIG on Agent-Mediated Electronic Commerce (AMEC), and four panel sessions for working groups.

After a brief welcome and some general remarks from the SIG coordinator Matthias Klusch, the morning session of the first day started off with a keynote talk of Norman Sadeh from the European Commission, DGXIII, about the pro-active initiative Universal Information Ecosystem (UIE) in the european research program IST (Information Society Technologies). He emphasised the intent of the UIE initiative to encourage in particular all kinds of 'revolutionary' projects towards R&D of new technologies used by highly dynamic population of infohabitants of emerging, adaptive, scalable and open information 'ecosystems' including future cyberspace (http://www.cordis.lu/ist/fetuie.htm). This talk attracted much interest and stimulated intense discussion on project ideas of attendees partially even late night the same day. The session continued with a block of four selected presentations on industrial and research collaborations as well as perspectives and concrete proposals for joint projects among SIG participants.

Vadim Ermolayev (Uni Zaporozhye, Ucraine) presented motives, perspectives and potential of R&D on intelligent agents for organisational management and teaching in eastern Europe, especially the Ucraine. Dieter Fensel (Uni Karlsruhe, Germany) contributed a short description of the ESPRIT project called IBROW3 on Internet-based brokering services for knowledge reuse including, for example, the development of ontology-mediated facilities to bridge heterogeneous information systems, an ontology editor, and a fairly plain graphical user interface of OntoBroker system. Chihab Hanachi (Uni Toulouse, France) reported on first steps towards a joint project with another SIG participant, the DAKE Center of the University of Keele, UK, on collaborative agent systems facilitating cross-organisational workflow among different companies/corporations. To end the presentation session, Klaus Fischer (DFKI GmbH) gave an outline of the so-called CASIMA project which deals with research and development of an agent-based pan-European intelligent information and trading network infrastructure for agriculture industry; the project is a collaborative effort of six AgentLink members.

In the afternoon, two panel sessions on the key topics of two working groups on agent-based mediation among information systems, data and knowledge management, and communication, coordination and collaboration (C3) among information agents were held and chaired by Dieter Fensel and Paolo Petta (Austrian AI Institute, Austria), respectively. The first panel started off with an overview from Sonia Bergamaschi (Uni Modena, Italy) about past and current R&D activities towards utilisation of coupled AI/Database techniques for integration of heterogeneous information sources including known approaches and systems like TSIMMIS, InfoSleuth, InfoMaster and MOMIS. Leon Sterling (Uni Melbourne, Australia) gave an account to XML-based integration of ontologies and sketched some of the work of his group on single, adaptive information agents for information retrieval in the Internet, like SportsFinder. Hans-Juergen Mueller (Deutsche Telekom AG, Germany) presented his group's work on composition and configuration platforms for Internet-based applications, in particular how to support effective knowledge management based on the concept of a so-called 'knowledge factory'. Hans Weigand (Uni Tilburg, Netherlands) shortly described the most prominent R&D projects carried out by his lab, like TREVI, MEMO and DECOMATE/CIA, in the areas of basic language ontology, electronic commerce, and use of description logic for intelligent integration of information, respectively. Finally, Richard Benjamins (Uni Amsterdam, Netherlands) briefly presented an approach for intelligent web page annotation and information search which was very lively debated. Due to the discussion and comments on given presentations the following three key topics and issues were identified for upcoming joint effort by SIG participants in that working area:

(1) Methods and standards for interoperability among heterogeneous, distributed information sources of different corporations,

(2) Intelligent information brokering and matchmaking; this includes, for example, agent capability description languages and tools, (re-)use of common, minimal or domain-specific ontologies and meta-data tools, and

(3) Agent-based support of (semi-automated) knowledge management in different enterprises and corporate virtual private networks or Intranets.

The second panel session was devoted to the working area of communication, coordination and collaboration (C3) among information agents. The panelists were Paolo Petta, Franco Zambonelli (Uni Modena, Italy), Luigi Serafini (IRST Trento, Italy) and Monica Divitini (Uni Trondheim, Norway). Coordination was seen as the sum of conventions and

commitments in a multi-agent environment, means as the process to ensure coherency (unity) of multi-agent systems. The presentations given by the panelists comprised, for example, the relation of C3 to CSCW in real-world context, the analytical distinction among basic C3 capabilities of an agent, the role of artefacts and possible approaches for metrics of C3 as well as the increasing need of learning and dynamic management of coordination in steadily changing collaborations of information agents. In general, the discussion mainly suffered from the fear of potential overlap with the objectives of the AgentLink SIG on Multiagent coordination and control. At the end several key topics and issues to be addressed by joint efforts uniquely in the I2A SIG have been worked out; the main topics are:

- Communication methods and standards, like intelligent, agent-based CORBA compliant service facilities for coordination and collaboration, basic shared ontology for C3, and advanced interfaces for standardised agent communication language (FIPA ACL),
- Formal coordination models and reference scenarios for collaborative information agents,
- Learning and metrics of quality of C3 in the context of cooperative information agents and systems.

The second day of the meeting started off with a short classification of intelligent information agents by the SIG coordinator who even briefly sketched some project proposals as one already accomplished goal of this meeting after the first day. The morning session was divided into one presentation from Martin Schneider (Siemens AG, Germany) describing the work of his group on a client/server-based agent broker for project matching, and two software demonstrations. The first demonstration was given by Pascal van Eck (Vrije Uni Amsterdam, Netherlands), standing in for his colleague Catholijn Jonker, on an intelligent Website architecture for information agents. Even interesting was the second live demonstration which addressed collaborative user profiling performed by the CASMIR agent system which has been developed at Uni Salford, UK.

The afternoon was devoted to the discussion of two more working groups, the third one focussing on the areas of Human-Agent interaction (HAI) and interfaces for information agents, the fourth one dealing with the relation of information agents and electronic commerce. The panel session on the third working group was chaired by Jeremy Pitt (Imperial College, UK), the panelists were Daniela D'Aloisi (FUB Rome, Italy), Mathias Bauer (DFKI GmbH, Germany), Vadim Ermolayev and Ian Dickinson (Hewlett-Packard Bristol, UK). After each panelist briefly presented the work of his group and even some ideas on the topics to be addressed by the working group, the following main issues and problems have been worked out in a discussion:

- Convenient inspection of agents by user, more transparency of agents' activity, clear impact of user feedback given in more native language, gestures or media input
- Standard of interface design for different kinds of information agents
- Need-driven not technology-lead products for agents in interfaces (which are not equal to complex agent architectures) and to cope with the danger of raised expectations of anthromorphisation.
- Shared context between user and intelligent interface agent including a common ontology for meaningful understanding in both directions, in particular to avoid that any single agent deployed on the Web will be just an individual curiosity to the user.

Since the last panel session intended to discuss one obvious key topic of the AgentLink SIG on agent-mediated electronic commerce, it was mutually agreed to held a common session of both SIGs. This common session was chaired by Walter van de Velde (Starlab Ltd., Belgium), the panelists were Markus Schwehm (Uni Stuttgart, Germany), Eugenio da Costa (Uni Porto, Portugal), Mikail Matskin (Uni Trondheim, Norway), Ygge (Uni Karlskrona/Ronneby, Sweden) and Klusch (DFKI GmbH, Germany). The discussion attracted much interest for several different reasons and was sometimes even a bit emotional concerning the evaluation of future vision and current state of the art of personalised agents doing everyday business on the Web including homebanking and online shopping for the average user. Some main questions and issues discussed by the panelists were, for example, how to increase the awareness of users regarding the potential of agent-based computing not only in business-to-customer electronic commerce, the mutual impact of electronic and material commerce, profit maximisation by online retailers vs. users' cravings for fun and social events, e.g., collaborative shopping (adaptive ShopBots for WebTV, etc.), notions of rationality not only in traditional terms of micro-economic based decision-making behaviour of trading agents. More detailed information on the outcome of this panel discussion will be available in a separate contribution by the panel chair. The status of a joint working group among AMEC and I2A SIG remains to be clarified. At the plenary meeting on April 23 the preliminary results of the meeting of every SIG and some additional news and administrative issues related to AgentLink as a whole have been reported.

Outcome of the Meeting

The second I2A SIG meeting had 43 registered participants, means a significant increase of interest since the kick-off meeting. Though, again twice as much attendees came from universities than from industry. But the large number of participants even showed the problem of satisfying different needs of active and passive attendees at the same time during a packed two-days meeting. However, the meeting was widely regarded as a successful event. In fact, much productive efforts towards collaboration and joint projects to propose for the recently started European IST research program have been put together, especially by fostering and establishing new inter-disciplinary contacts among attendees of this meeting in London.

The core activities of the two-day meeting were strongly determined by the panel sessions on the working groups intended to focus the work of the I2A SIG on a few important topics and to produce significant results theoretically and in practice in reasonable time. Actually, three important sub-areas of information agent technology have been discussed for future joint efforts among SIG participants, thereby forming kind of domain-specific working groups, at the meeting. Those groups focus on

- Agent-Based mediation among information systems, data and knowledge management (WG-1)
- Communication, coordination and collaboration among information agents (WG-2)

• Human-Agent interaction and interfaces for information agents (WG-3)

In mid-term it is expected that each of these working groups will participate in national and European research programs, especially by joint projects among SIG participants in the respective working areas. First steps were taken already to accomplish this objective.

Future work

To further improve the productivity of upcoming I2A SIG meetings it is intended, for example, to restrict the number of participants to a much lesser amount, and to try to 'recruit' more industrial attendees. Besides, discussion on current or potential joint projects within the SIG working groups, means collaborative work and efforts among SIG participants in the mentioned working areas, will be the dominant factor of the next meeting. The next I2A SIG meeting will be held in Barcelona, Spain, in late September this year (probably September 20 - 22, 1999).

The technological roadmap of the I2A SIG for the forthcoming period including that of each of its three working groups will be available as a kind of comprehensive green paper in late summer. In addition, it is intended to publish the preliminary results from substantial efforts within the I2A SIG as a journal article or book at the end of this year.

To facilitate convenient communication among SIG participants an open public majordomo mailing list infoagents@gmd.de, which anyone interested in is invited to submit contributions, as well as a (protected) shared working space (http://bscw.gmd.de/bscw) for exchange of more sensitive documents, like pre-proposals on joint projects among SIG participants, have been created and maintained. Several contacts to related research communities, like that for database and information systems, and information retrieval have been established.

More information on this SIG can be found at:

http://www.dfki.de/~klusch/i2a-SIG.html

Multi-agent Coordination and Control

John Perram
The Maersk Mc-Kinney Moller Institute for Production Technology, SDU-Odense University, Denmark
jperram@mip.sdu.dk

Multi-agent concepts and methodologies are finding increasing application in controlling complex, unpredictable systems in real time. Agent-like approaches routinely accept uncertainty and distribution, leading to control schemes where decision-making and responsibility are distributed much more widely than in conventional engineering practice. The successes that they have begun to produce are likely to lead to their heavier use, including use in applications that have time-critical aspects. In order for them to succeed in such situations, however, they will need development that addresses real-time and time-critical properties explicitly.

The subdivisions following the one immediately below are labelled in terms of applications, because at present we consider it desirable to drive the selection of basic-research topics by reference to concrete problems that are turning up in specific applications. This SIG bring together researchers from a number of areas fitting into the paradigm, including:

- Adaptive real-time systems with knowledge-based components
- Manufacturing systems
- Robotics including multiple robot systems and robot football.

Adaptive real-time systems with knowledge based components includes traffic engineering applications. Multiple robot systems could include collections of robots which cooperate to perform some form of useful activity such as assembly operations or playing football or other games. The area of intelligent manufacturing systems aims to use multi-agent systems both as a modelling tool and a control software system.

The challenge for the SIG is to identify and focus on the general issues of coordination and control which are important for such applications, and to use the applications to stimulate and test advances in research on the issues, for example:

- How and where to represent information about the systems and the agent groups that operate on them
- Mechanisms (e.g. varieties of caching) for delivery of adequate responses to time-critical demands
- Coordination methods that are computationally economical enough for real-time use
- Decentralised management of limited common resources
- Adaptation of results from machine learning to collective learning in multi-agent systems
- New schemes of decentralised control to take account of the "real time" dimension.

The "coordination and control" emphasis in this SIG is because the development of adequate methods for coordination and for control are crucial for the successful operation of multi-agent systems when a real-time discipline is imposed. Progress in other multi-agent topics is desirable too for real-time applications, but it will not be particularly helpful unless the problems of coordination and control are resolved adequately.

The first meeting

About 20 members and associated members attended the first SIG meeting held at Imperial College, London in April, and set themselves the task of trying to

extract a framework for coordination and control by studying the various application areas which were represented. A second task of creating a technological roadmap for the network was also addressed. Participants were roughly equally divided between robot football, multiple robot systems, intelligent manufacturing systems and traffic engineering. In addition, there were a few participants from the world of software engineering to help with the process of distilling a framework for applications in which spatio-temporal coordination was an important issue.

Each group had chosen a "spokesman" who presented the coordination issues from their respective domains. Bo Jorgensen provided an overview of coordination seen from a software engineering point of view. An important, common, coordination mechanism seemed to be the need to satisfy constraints imposed by the application, such as requiring that robots don't simultaneously occupy the same physical space. The general feeling was that considerable progress towards the first goal had been made. It was proposed to construct a set of questions to be considered by participants for the next meeting, to be held in Barcelona during 5-7 July, immediately after MAAMAW, which will take place in Valencia. This early date was chosen because we felt that we had achieved a certain amount of momentum and synergy.

We also spent some time discussing the roadmap, which should provide answers to the following questions:

- Where do we want to be (i.e. what is our vision)?
- Where are we now?
- How do we get from here to there?

The vision which emerged was of how the next phase of the IT revolution, in which intelligence will be built into to mechatronic components, could benefit from multi-agent concepts. We foresaw this leading to a customized society in which people, software and machines would all be first-class citizens. Developments such as the modular Mindstorm components being produced by the toy manufacturer Lego, represented at the SIG by the LEGOlab from the University of Aarhus, were seen to be significant. Such distributed hardware systems as these, robot football teams, autonomous vehicle control systems and distributed manufacturing systems, were also seen as important testbeds for researching the nature of agency in systems with an important spatio-temporal element.

The PowerPoint slides of the coordinator's plenary presentation can either be obtained by requesting them from him or from the SIG home page.

It was also felt that it is difficult to partition the field of MAS into independent SIG's, so it is hoped that any other AgentLink members who feel that they belong with us should contact the coordinator.

More information on this SIG can be found at:

http://www.AgentLink.org/activities/sigs/sig5.html

Agent R&D in Europe

Mike Wooldridge
AgentLink, Queen Mary and Westfield College, UK
M.J.Wooldridge@elec.qmw.ac.uk

AgentLink is now a year old, and in that year, the network has grown from 37 members to over 100. Given that we have now passed the magic 100, it seemed appropriate to ask ourselves exactly who our members are and where they come from. The results should be taken with a pinch of salt – 100 members hardly makes for a statistically significant survey. But they do at least give an indication of who is doing agent R&D, and where this is happening.

The most obvious analysis we can do is to look at number of members of AgentLink by country – see Figure 1. The UK and Germany clearly have the strongest showing; the UK figure is particularly high. Perhaps the most encouraging single observation we can make is that Europe is now well covered by AgentLink – we have representation and activity throughout the region.

Perhaps more enlightening is to look at the number of AgentLink nodes per million of population. It turns out that, on average, there are 0.26 AgentLink nodes per million of population throughout Europe. Figure 2 shows the normalised

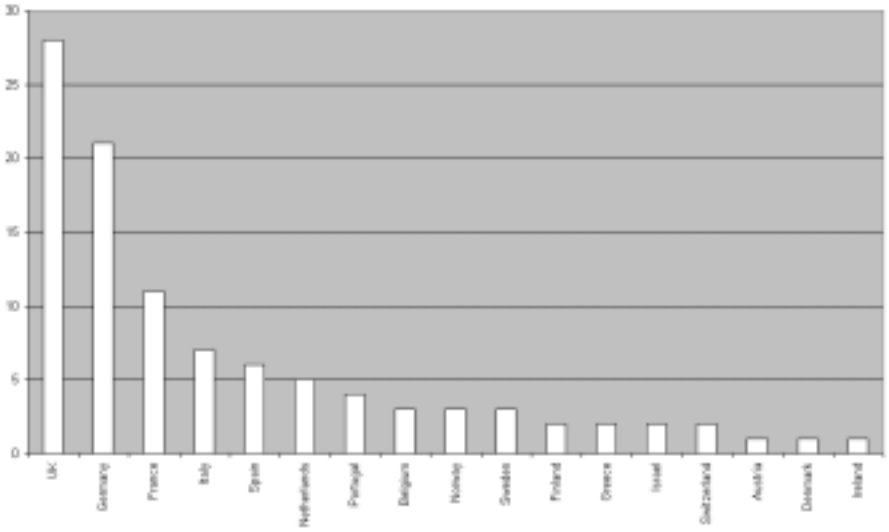

Figure 1: AgentLink membership by country

value. Germany represents almost exactly the average; France, Italy, Spain, Greece, Austria, and Denmark are all below average, while the UK, Portugal, Norway, and Finland all have above average representation. The deviations are not great, however; most countries seem to have a representation approximately on par with their population size.

We also examined the membership of AgentLink by organisation type (university, industry, research institute, or public administration). The results are shown in Figure 3. Perhaps surprisingly, AgentLink has only 55% university membership. Industrial and public administration together account for about 30%, with another 20% from research institutes. Our goal was to aim for no more than

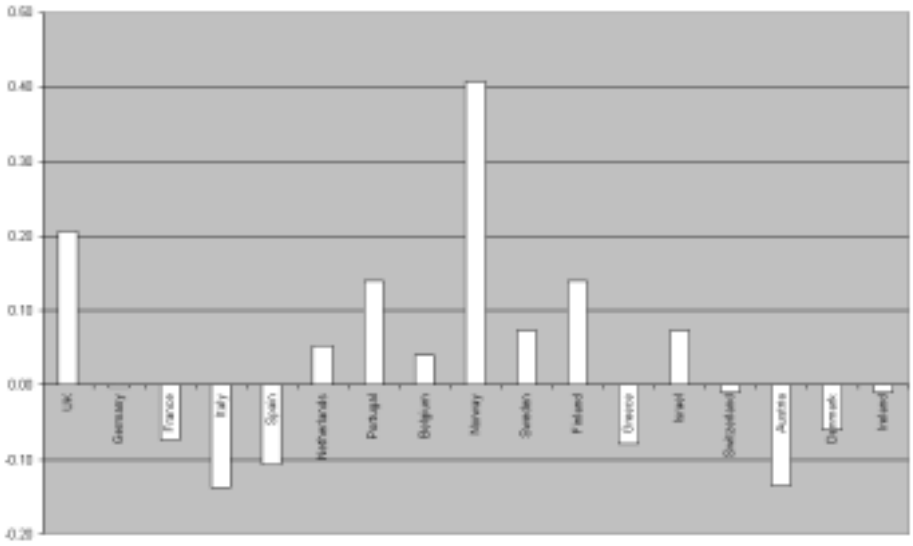

Figure 2: AgentLink membership by country, normalised by population size

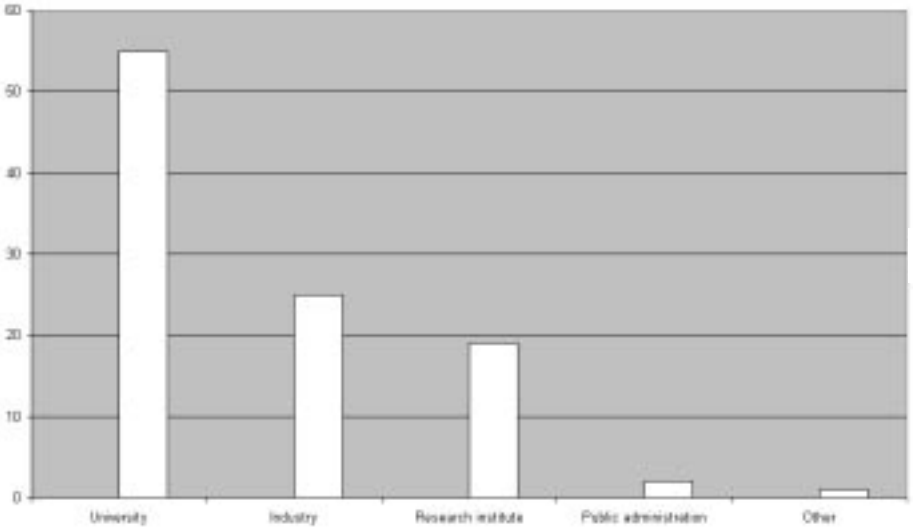

Figure 3: AgentLink membership by organisation type

The next phase of AgentLink recruitment will focus primarily on industrial members. There are still many potential industrial participants who either do not know about AgentLink, or else do not know about agent technology. Accordingly, a major thrust of the industrial action component of AgentLink over the coming months will involve raising awareness with such potential members. Although we have made a strong start (nearly tripling membership in one year is excellent), there is still a long way to go in ensuring that the community at large really do understand how AgentLink can serve them.

60% academic membership, with 40% industrial membership, so these figures are encouraging.

Finally, we analysed industrial members by their industry area. (We took the information from registration forms completed by members when they joined the network.) The results are shown in Figure 4. Not surprisingly, software has a strong showing, but the two telecoms categories are together almost as strong. What is most perhaps most interesting is the *spread* of members, across industrial sectors. It is encouraging to see pharmaceutical companies expressing an interest in agent technology, for example.

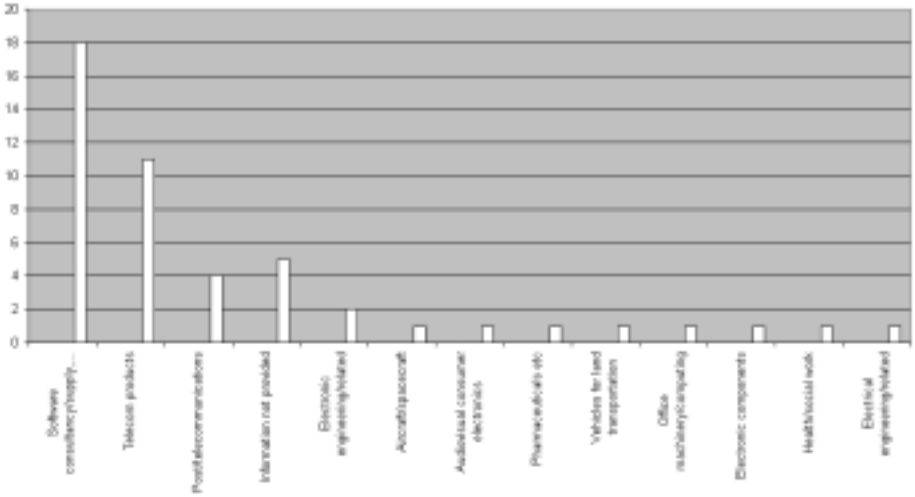

Figure 4: AgentLink membership by industry area

2000		
IMACS 2000	Special Session on Agent-Based Simulation, Planning and Control at 16th IMACS World Congress 2000, Lausanne, Switzerland.	21-25 August 2000 http://www.sztaki.hu/~vancza/imacs/agents.html
ICMAS 2000	Fourth International Conference on Multiagent Systems Boston, USA.	July 7-12, 2000 http://www.lania.mx/icmas/
Agents 2000	Fourth International Conference on Autonomous Agents Barcelona, Spain.	June 3-7 2000 http://www.iiia.csic.es/agents2000/
SAC 2000	2000 ACM Symposium on Applied Computing Como, Italy.	March 19-21, 2000 http://www.cs.ucy.ac.cy/SAC2000.html
HICSS-33	Minitrack on MultiAgent Systems, Internet & Applications at 33rd Hawaii International Conference on System Sciences, Hawaii, USA.	January 4-7, 2000 http://cs.unomaha.edu/~rewini/SWT-CFP.html

1999		
IAT'99	1st Asia-Pacific Conference on Intelligent Agent Technology Hong Kong.	December 14-17, 1999 http://www.comp.hkbu.edu.hk/IAT99
SCAI'99	Scandinavian AI Conference'99 Odense, Denmark.	December 12-15, 1999 http://www.mip.sdu.dk/~scai99/first.htm
MAS'99	Multi-Agent Systems in Production Vienna, Austria.	December 2-4, 1999 http://www.ihrt.tuwien.ac.at/MAS99/
ICLP'99	Workshop on Multi-Agent Sytems in Logic Programming held in conjunction with ICLP'99, Las Cruces, New Mexico, USA.	December 4, 1999 http://www.cs.nmsu.edu/~complog/conferences/iclp99/
EC-99	ACM Conference on Electronic Commerce Denver, Colorado, USA.	November 3-5 1999
MATA'99	First International Workshop on Mobile Agents for Telecommunication Applications, Ottawa, Canada.	October 6-8, 1999. http://ocri.genie.uottawa.ca/mata99/
ASAP'99	The First International Symposium on Agent Systems and Applications, California, USA.	26-29 Sep http://www.generalmagic.com/asap
IMS'99	Intelligent Manufacturing Systems 1999 Katholieke Universiteit Leuven, Belgium	September 22-24, 1999 http://www.mech.kuleuven.ac.be/pma/project/imswg/ims99
Virtual Agents 99	One day workshop prior to the UK Virtual Reality SIG Conference Salford, UK.	September 13, 1999 http://www.salford.ac.uk/cve/va99/
ACTS Workshop	Advanced Services in Fixed and Mobile Telecommunications Networks Singapore.	9-10 September 1999 http://www.comnets.rwth-aachen.de/~cameleon/acts_workshop.html
DOA'99	International Symposium on Distributed Objects and Applications Edinburgh, Scotland.	5-6 September, 1999 http://www.cs.rmit.edu.au/conf/doa99
IATA'99	3rd International Workshop on Intelligent Agents for Telecommunication Applications, Stockholm, Sweden.	9-10 August, 1999 http://dai.cs.tu-berlin.de/english/news/tagungen/
IJCAI'99	Sixteenth International Joint Conference on Artificial Intelligence Stockholm, Sweden.	13 July-6 Aug http://www.dsv.su.se/ijcai-99/
CIA-99	Third International Workshop, Cooperative Information Agents, Stockholm, Sweden.	31 Jul-2 Aug http://www.informatik.tu-chemnitz.de~klusch/cia.html
ATAL'99	Agent Theories, Architectures and Languages - Sixth International Workshop, Orlando, Florida, USA	July 15-17, 1999 http://www.elec.qmw.ac.uk/dai/atal

About *AgentLink News*

The aim of the AgentLink newsletter is to provide a relatively informal way of communicating both what's happening in AgentLink, but also what's happening in the agent world generally. Many newsletters end up being rather dull. (Let's face it, the very name "newsletter" puts many people off.) AgentLink News aims to be different. In addition to containing the worthy-but-dull details of what's happening in the network, we aim to carry a range of articles including features, reports on conferences and workshops, informal descriptions of research results and new software, book reviews, and so on. Of course, we can't do this without your help! We need you to generate the content for the newsletter. If you are interested in writing something for the newsletter, please get in touch with the editor Paul Davidsson directly, at the address below. The deadline for receipt of articles for issue three is September 3, 1999. Remember: AgentLink news is not an academic journal, so we won't publish academic articles. Pieces should follow the conventions of similar sorts of publications (such as AI Magazine, or IEEE Internet), and be relaxed in style, with short lists of references, etc. Length is also an issue – features should be no more than a few pages.

The newsletter will have a circulation of up to several thousand. It's an ideal way to communicate your work to a specialist community, who will want to hear about it. So why not contribute?

How to Contact AgentLink

AgentLink general coordinator
Michael Wooldridge
Department of Electronic Engineering
Queen Mary and Westfield College
University of London, London E1 4NS
United Kingdom
Email: coordinator@agentlink.org

AgentLink administrator & WWW manager
Hugo Brailsford
Department of Electronic Engineering
Queen Mary and Westfield College
University of London, London E1 4NS
United Kingdom
Email: coordinator@agentlink.org

Industrial Action (workpackage 1) coordinator
Joerg P. Mueller
Siemens AG
ZT IK 6
D-81730 Munich
Germany
Email: joerg.mueller@mchp.siemens.de

Research (workpackage 2) coordinator
Yves Demazeau
Laboratoire LEIBNIZ - Institut IMAG
46, avenue Felix Viallet
38000 Grenoble
France
Email: Yves.Demazeau@imag.fr

Teaching & Training (workpackage 3) coordinator
Gerhard Weiss
Institut fuer Informatik
Technisch Universitaet Muenchen
D-80290 Munich
Germany
Email: weissg@informatik.tu-muenchen.de

AgentLink News Editor
Paul Davidsson
Department of Computer Science
University of Karlskrona/Ronneby
372 25 Ronneby
Sweden
Email: Paul.Davidsson@ide.hk-r.se

gentLink News design by Hugo Brailsford