European Coordination Action for Agent-based Computing

AgentLink News

ISSUE 15 | September 2004 | ISSN 1465-3842

The end of summer heralds the return of many things: of students to their Autumn classes, of holiday-makers back from the beach, and, of course, a feature-filled AgentLink III Newsletter, currently in its 15th issue. This issue has a semantic flavour: in addition to the regular AgentLink articles, we have brought you the first of a two part feature on the EU FP6 "Semantic-Based Knowledge Systems" Network of Excellence Projects. Three Networks are featured in this issue: REWERSE, MUSCLE, and Knowledge Web, with two further Networks to appear in the next issue of the AgentLink Newsletter.

This issue opens with a feature article by Ian Dickinson (HP Labs Bristol, UK), which explores the synergy between Agents and the Semantic Web. Continuing with the semantic theme, Beun and van Eijk highlight the problems that may arise due to ontological discrepancies that can occur between different agents (a potentially significant problem when several agents from different sources/communities collaborate, such as within the Agentcities framework), and present an approach for tackling this problem. Following on from last issue's report by Steven Willmott on current activities of the Agentcities/openNet Testbed, Julian Padget reports on the AgentcitiesUK Network, and the upcoming Challenge Day, to be held at the University of Bath. Agent-based research has profited from analogies in various fields, such as economics and biology. In "Socially Inspired Engineering", Hales and Edmonds discuss the use of results that originated in agent-based simulation of social behaviour for solving open problems in the engineering of multi-agent systems. Agent societies, or institutions, also feature in Ulises Cortes's article on "Electronic Institutions and Agents", whereby five recent PhD theses in this area, presented throughout Europe, are reviewed.

An important task within the AgentLink portfolio is to support and sponsor Agent-related events across Europe, and to report on these activities. Reports of six recent conferences and workshops are presented in this issue, showing the diversity of current activity across the Agent Community. Likewise, standards activities continue to progress, and Monique Calisti highlights some of the more significant activities in her regular Standards column. Completing this issue, we have a review by Terry Payne of Lin Padgham and Michael Winikoff's recently published book on "Developing Intelligent Agent Systems", and an industrial report on the Cybele Agent Infrastructure.

Although the summer may be drawing to an end, activity within the AgentLink community continues to grow, with many AgentLink III sponsored conferences and events scheduled for the coming months, as well as the forthcoming Agent Technology Conference to be held in Zurich in October 2004. With the 2nd AgentLink Technical Forum planned for early 2005, the liveliness of next few months for the Agent research community only reassures its ever expanding role in IT research and development across Europe!

Terry Payne
AgentLink Coordinator
University of Southampton

Rafael Bordini
AgentLink News Editor
University of Durham

CONTENTS

The Semantic Web and Software Agents: Partners, or Just Neighbours?

Ian Dickinson
HP Labs Bristol
United Kingdom
ian.dickinson@hp.com

A stated goal of the Semantic Web is to make web-based information accessible to machine processing, where currently it requires human mediation [1]. One view of agent technology is precisely that it is concerned with creating autonomous computational processes that mimic, or augment, human behaviour. On the face of it, these would seem very complementary technologies, producing many examples of agent systems that exploit Semantic Web information sources. While there are some good examples (e.g. [2]), these technologies are not used together as often as might be expected. In this brief article I will try to sketch some of the opportunities for using the Semantic Web with agents. I will also note some challenges that both research communities should address.

What is the Semantic Web?

A common, but unfortunate characteristic of both agent technology and the Semantic Web is that they resist precise definitions. Both technologies share a common root desire, however. We want computers to do more things for us, particularly things that currently require human knowledge and expertise. Booking travel, monitoring sensor networks, trading stocks and shares, and recommending music, films and restaurants are among the many examples that have been suggested. In essence, the idea behind the Semantic Web is very simple: by changing the information structures on the web, making them more explicit and meaningful, we can begin to develop rich, interconnected knowledge sources. These knowledge sources can then support the automation goals we seek.

Currently, information is overwhelmingly presented on the web under the assumption that it will be read by people via a web browser that renders HTML, but HTML is a presentation rather than information encoding language. As anyone who has ever had to write a web-scraping script knows, almost any amount of logical confusion in the HTML code can be tolerated (and found in published pages) as long as the end result looks OK when people read pages via browsers. Automated processes, whether agents or other programs, cannot rely on the visual presentation of information to determine logical structure. A shopping agent is likely to be confounded if the only association between a price and product code is that they appear on consecutive rows of a layout table.

If this sounds familiar, it should. Similar things — the separation of style and content — were said of the introduction of XML. XML, however, is basically just a syntax for data representation. The Semantic Web attempts to incrementally build on both the WWW and on XML, by adding a formalism for expressing the meaning of the encoded information. It should be noted, however, that the aim is evolution of the current WWW, not a *revolution* that will replace it.

Although the Semantic Web is a general idea, the term is closely associated with a set of technology standards from the W3C. These build a stack of specifications for increasingly sophisticated languages, from RDF to OWL (see the "layer-cake" diagram in [1]). Other specifications, such as a standard query language, and a set of best-practice guidelines, are still under development by W3C working groups.

RDF and OWL — a capsule summary

I will only summarise here the standard notations defined by the W3C for basic data representation on the Semantic Web. The *Resource Description Framework (RDF)* [3] defines a simple model in which data is encoded as a directed, labelled graph. The edges in the graph are labelled with a predicate symbol, so the graph can also be thought of as a set of 3-tuples, or *triples*:

subject-node — predicate symbol — object-node

In RDF, nodes are either *resources* (i.e. symbols), or data *literals*. A resource is typically identified by a *Uniform Resource Identifier* (URI). A common URI type is the HTTP URL (e.g., http://www.hp.com), but there are many other types. A literal is a concrete data value, such as an integer or string. One special resource type is the *unlabelled* or *anonymous resource*. An anonymous resource has the blank label, so it is often referred to as a blank node or bNode. I will omit discussion of the RDF schema language (RDFS); there is extensive reference material on the W3C site [3] and elsewhere.

OWL is the ontology language in the W3C specification stack. OWL extends RDF by specifying terminology for declaring axioms that define classes, class expressions and properties in a richer *description logic*. Description logic is a subset of the predicate calculus. The choice of language constructs is made so as to ensure the logic is *decidable*. In OWL, classes may be named, or they may be defined as Boolean expressions of other classes: the *union* operator is equivalent to logical disjunction, *intersection* to conjunction, and *complement* to negation. The sub-class re-

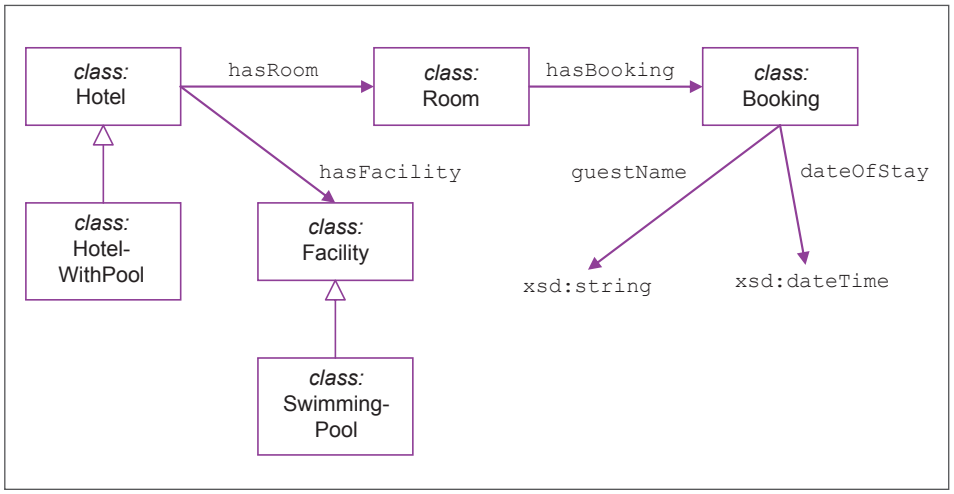

Figure 1: A Simple Ontology for the Hotel-booking Domain

```
<?xml version='1.0'?>
<!DOCTYPE rdf:RDF [
    <!ENTITY rdf   'http://www.w3.org/1999/02/22-rdf-syntax-ns'>
    <!ENTITY rdfs  'http://www.w3.org/2000/01/rdf-schema'>
    <!ENTITY owl   'http://www.w3.org/2002/07/owl'>
    <!ENTITY hotel 'http://example.org/hotel'>
]>
<rdf:RDF
   xmlns:rdf="&rdf;#" xmlns:rdfs="&rdfs;#" xmlns:owl="&owl;#"
   xml:base="&hotel;">
  <owl:Ontology rdf:about="" />

  <owl:Class rdf:ID="Booking" />
  <owl:Class rdf:ID="Hotel" />
  <owl:Class rdf:ID="Room" />
  <owl:Class rdf:ID="Facility" />
  <owl:Class rdf:ID="SwimmingPool">
    <rdfs:subClassOf rdf:resource="#Facility" />
  </owl:Class>

  <owl:Class rdf:ID="HotelWithPool">
    <owl:equivalentClass>
      <owl:Restriction>
        <owl:onProperty rdf:resource="#hasFacility" />
        <owl:someValuesFrom rdf:resource="#SwimmingPool" />
      </owl:Restriction>
    </owl:equivalentClass>
  </owl:Class>

  <owl:ObjectProperty rdf:ID="hasRoom" />
  <owl:ObjectProperty rdf:ID="hasFacility" />
  <owl:ObjectProperty rdf:ID="hasBooking"/>

  <owl:DatatypeProperty rdf:ID="dateOfStay" />
  <owl:DatatypeProperty rdf:ID="guestName" />
</rdf:RDF>
```

Figure 2: Simple Hotel Ontology Represented in OWL/XML

lationship on classes expresses implication, while quantifiers are expressed by *restrictions*, which constrain a property to have *all values from* a given class (universal quantification) or *some values from* a given class (existential quantification). Property axioms define a predicate hierarchy, and allow predicates to be denoted as *transitive*, *symmetric*, *functional* or *inverse functional*. These capabilities taken in full, together with the semantics of RDF, are in fact not decidable. To preserve the decidability of description logic, a set of syntactic rules restrict the use of some of OWL Full's constructs to form what is termed OWL-DL. Details of the representational capabilities of OWL, and the rules that define OWL-DL, are available from the W3C site [4].

Of syntax. The most common complaint about RDF is that it is ugly XML. Other surface syntaxes for RDF are more friendly for human readers, but are equivalent in meaning because the semantics of RDF and OWL is defined in terms of the underlying graph and information model, not the serialisation syntax.

A simple OWL example

To illustrate some of these ideas, Figure 1 shows a very simple hotel ontology, while Figure 2 shows one possible OWL encoding of it.

The declaration of `HotelWithPool` shows the use of `owl:someValuesFrom`, the analogue in description logic to existential quantification. The class `HotelWithPool` is exactly equivalent to the set of things that have one or more facilities (the `hasFacility` property) of type `SwimmingPool`. Resources of type `HotelWithPool` can have other facilities, but they must have at least one pool. The remaining lines name the properties that relate the resource nodes to each other. In OWL, *object properties* have other OWL resources in their range, whereas *datatype properties* have data literals as their range values.

1. Aduna BV. Sesame RDF platform. 2004. (http://www.openrdf.org/)
2. Mindswap Lab, U. o. M. Pellet OWL Reasoner. 2004. (http://www.mindswap.org/2003/pellet/index.shtml)
3. Tucana Technologies. Kowari Metastore. (http://kowari.sourceforge.net/)
4. Dave Beckett. Redland RDF Application Framework. 2004. (http://www.redland.opensource.ac.uk/)
5. Stanford University. The Protégé Project. 2004. (http://protege.stanford.edu/)
6. Kalyanpur, A. SWOOP. 2004. (Semantic Web Ontology Overview and Perusal) (http://www.mindswap.org/2004/SWOOP/)

Semantic Web Tools

The availability of software tools for processing Semantic Web information sources is increasing steadily. Java programmers are particularly well served. Available Java-based platforms include Jena [5], Sesame[1], Pellet[2] and Kowari[3], all of which are free for download and most are open-source. In addition, there are RDF toolkits for C/C++[4], PHP, and JavaScript. These platforms allow programmers to read, parse, manipulate, query, and persist RDF and OWL information via a convenient API. Many of the tools perform inference using the underlying semantics of the languages. Semantic Web documents can be edited in source code form (e.g., RDF/XML), but there are numerous visual editing tools available, for example, Protégé[5] and SWOOP[6].

Agent applications with RDF and OWL

As discussed above, we have a set of specifications for Semantic Web languages, which are designed to allow developers to move towards the grand vision by building upon and extending the current web. So far, so good — but where do agents fit in?

Practical uses of the Semantic Web with agents

Below, I outline four generic ways in which today's Semantic Web technology can be of direct benefit to agent designers. I will address four topics: controlled vocabularies, semi-structured data, query, and reasoning.

1. Creating controlled vocabularies

When developing applications that use symbolic computing, it is important that the symbols correspond wherever they are used. If one developer uses `hotel:guestName`, while another uses `hotel:guest_name`, bugs are inevitable. A *controlled vocabulary* is a set of symbols (class names, predicate names, etc.) that are permissible for a given application. Simply writing the domain ontology as an OWL document helps to define the controlled vocabulary, but it is then up to individual developers to check that they are using the symbols from the controlled vocabulary correctly. A better way to ensure symbols are used correctly is to define a set of program constants which developers can then use directly in their code. In Java, we might define the constant `GUEST_NAME` in the `HotelVocab` class as a static final variable. However, this suffers the problem that the constant declarations in Java and the OWL file may get out-of-sync, again introducing bugs. Tools such as Jena's *schemagen* [5] address this problem. Schemagen automates the process of creating constants from a set of OWL or RDFS declarations, so a developer team needs to maintain only one source file, the OWL ontology. Generation of the constants in the source code is then an automated step in the project's build process.

2. Using semi-structured data

RDF applications use representations built from

```
select ?r where (?r rdf:type hotel:Room)
                (?r hotel:hasBooking ?b)
                (?b hotel:dateOfStay
                    "2004-07-30T00:00:00.000-00:00"^^xsd:dateTime)
using hotel for <http://example.org/hotel#>
```

Figure 3: Example RDQL Query

graphs of linked nodes. This encoding is sometimes termed *semi-structured*, because it has a definite structure, unlike natural language text, but it is not as rigidly fixed as, for example, a relational database schema. In a semi-structured dataset, if you need to add in a new node to encode some new condition in the data or problem, you can. As long as the data-processing agent is able to interpret the meaning of your node and predicate labels, this flexibility can be freely exploited. XML data is also semi-structured, but the semantic model underlying RDF and OWL give some hope that the data-processing agent can make use of data that it has not necessarily seen before. Indeed, a key strength of RDF over XML is the open nature of RDF datasets: additional data can be linked to a resource without needing to side-effect the original document. This additional information can be more meaningful in RDF or OWL than it is possible with XML schema languages.

A FIPA ACL message can use RDF directly as a content language. Requesting a receiving agent to perform some action is a typical case where a semi-structured data formalism is useful. The primary verb (the message performative) is `request`. The message content must then convey both the identity of the action to perform, and any parameters it takes. The action could be an RDF resource node, with attached properties indicating the action type and arguments. Agents can share ontologies of actions types and argument properties that are published within a given agent community.

3. Using RDF queries

There are a number of existing query languages specialised for querying RDF knowledge sources. Efforts are underway to define a standard query language [6]. Until the new language is finalised, tools such as Jena support RDQL [7], a widely deployed RDF query language. RDQL uses an SQL-like syntax, specialised for traversing graphs of RDF triples. An example query in the hotel booking domain is shown in Figure 3.

Handling such a query with an agent platform raises some interesting design choices. A natural place to start might be the query communicative acts. However, neither of the standard communicative acts in FIPA00037 quite fit: `query-if` expects a predicate, requiring the receiver to respond with either true or false (or some equivalent encoding), while `query-ref` takes an identifying reference expression. An RDQL query, or indeed a query encoded in another query language, does not conform to the pattern for an identifying reference expression. This seems to leave `request` as the only choice from the standard FIPA communicative acts, assuming that the `select` verb that introduces an RDQL query can be viewed formally as an action name. Clearly, any given agent community can, by design, adopt a new performative with the meaning "execute this query and return the results". However, this then closes the community to participation by otherwise standards-conforming agents.

> ...if the Semantic Web delivers on the promise of providing large-scale richly structured information sources, then many agent designers may well find significant benefits from adding Semantic Web technology to their applications.

```
(request
  :sender a
  :receiver b
  :language "http://jena.hpl.hp.com/2003/07/query/"
  :content "select ?r where (?r rdf:type hotel:Room)
                            (?h rdf:type althotel:PoshHotel)
                            (?h hotel:hasRoom ?r)
            using hotel    for <http://example.org/hotel#>
                  althotel for <http://example.org/alt-hotel#>"
  :ontology "http://example.org/alt-hotel" )
```

Figure 4: Referencing an Alternative Definition

4. Using OWL inference

I conclude this brief trip around some of the points of contact between agents and the Semantic Web by looking at an inference example. Suppose that a new agent arrives in the agent community, with an alternative to the standard ontology. For example, this new arrival defines a *posh hotel* as one with a swimming pool. This declaration is declared to be a sub-class of the existing class `hotel:HotelWithPool`. Figure 4 shows a query message using terms from this new ontology. Note that I have explicitly listed the URI of the alternative ontology in the `:ontology` field of the ACL message.

One way for the receiving agent to deal with this situation is to construct, on demand, a new knowledge base that contains the current database of assertions about bookings together with the original ontology and the new definitions provided by the requesting agent. In Jena, for example, this is done with a few lines of Java code to build a new *inference model* and to add the new definitions as imported sub-models. Once this new model has been constructed, it can be queried using the RDQL query directly from the message. The inference model uses the OWL semantics to find entailments from the ontologies and the prior assertions so as to determine if there are any known posh hotels with bookings.

Issues and challenges

Briefly, here are a few of the practical and theoretical issues that need to be addressed to make the best use of both agents and semantic web technologies.

Open world assumption

Attentive readers will have noticed that the expression in Figure 4 queries for a booking — that is, a room that has a booking record attached. Perhaps more useful would be a query for a free room. However, there is no negation in RDQL, so a query that fetches rooms that do *not* have a booking cannot be defined in RDQL. One reason that RDQL has no `NOT` operator is that the Semantic Web, in general, makes the *open world assumption*. Under the open world assumption, lack of information about some proposition does not mean that the proposition is assumed false. Just because a given agent cannot locally

see a statement that room x is booked does not mean that there is no such statement somewhere "out there" in the information space. Under the closed-world assumption, an agent would be free to presume that a room that was not known to be booked was free.

This open world assumption often misleads developers who are more accustomed to working in closed world domains. However, if the Semantic Web is truly to deal with web-scale information sources, there is no really effective alternative. The challenge, then, is for agent designers to work with the open-world semantics in mind, or to consciously layer a private closed-world semantics on top a standard open-world semantic-web infrastructure.

Standards mismatch

Agent standards, such as FIPA, matured before the widespread use of RDF and similar technologies. This has resulted in a number of potential incompatibilities. For example, symbols in the Semantic Web are URI's, either written out in full or abbreviated in q-name form (i.e., `prefix:name`). URI's are not strings, but must be written as strings in FIPA ACL for consistency with the ACL grammar. Some performatives are missing (such as `query`; see above). On the Semantic Web, ontologies are typically referenced by URL, and a given fragment of OWL may reference a number of ontologies by name. This makes the use of the `:ontology` message parameter unclear. None of these problems are insurmountable, but if every agent project adopts different conventions, interoperability is reduced.

Web services and Semantic Web services

Perhaps a more fundamental question is the overlap between agent infrastructure and web services [8]. Many of the services of a basic agent society — message passing, directory services, security, message brokering, etc. — are also available through web-services platforms. With the development of Semantic Web services (e.g., web service descriptions using OWL) [9], traditional agent capabilities such as delegation, collaboration, and federation are being adopted into the web services model. Moreover, web services are apparently already more widely deployed than agent infrastructures, and arguably have a more active development community. At the least, we can expect to see gateways that link the worlds of agents and web services. Perhaps, in the future, agents will rely more on web-services infrastructure for their middleware than on agent platforms *per se*.

I have outlined some of the ways in which Semantic Web technology and agent technology can complement each other, and touched on a few of the open issues that still need to be addressed. There is much more exploration to be done in the combination of both technologies, but if the Semantic Web delivers on the promise of providing large-scale richly structured information sources, then many agent designers may well find significant benefits from adding Semantic Web technology to their applications. At HP, we aim to continue to contribute to this goal through applied research and supported, open-source, developer tools and platforms such as Jena.

Acknowledgements

Thanks to Jeremy Carroll, Chris Preist and Dave Reynolds of HP Labs for their comments on an earlier draft of this article.

References

1. Berners-Lee, Tim, Hendler, James, and Lassila, Ora "The Semantic Web". *Scientific American*. 2001.
2. Grimnes, G. A., Chalmers, S., Edwards, P., & Preece, A. "GraniteNights — A Multi-Agent Visit Scheduler Utilising Semantic Web Technology". In: *Proc. Seventh International Workshop on Cooperative Information Agents (CIA-2003)*. 2003. pp. 137-151.
3. World Wide Web Consortium (W3C). *The Resource Description Framework (RDF)*. 2004. http://www.w3.org/RDF/
4. Bechhofer, Sean, van Harmelen, Frank, Hendler, Jim, Horrocks, Ian, McGuinness, Deborah L., Patel-Schneider, Peter F., and Stein, Lynn Andrea. *OWL Web Ontology Language Reference*. W3C. 2004 . http://www.w3.org/TR/2004/REC-owl-ref-20040210/
5. HP Laboratories. *The Jena Semantic Web Platform, version 2. 2004.* http://jena.sourceforge.net
6. W3C. *RDF Data Access Working Group*. 2004. http://www.w3.org/2001/sw/DataAccess/
7. Seaborne, A. RDQL — *a query language for RDF. 2004.* http://www.w3.org/Submission/2004/SUBM-RDQL-20040109/
8. Alonso G, Casati F, Kuno H, Machiraju V. *Web Services: concepts, architectures and applications*. Springer, 2004.
9. OWL-S coalition. *DAML services*. 2004. http://www.daml.org/services/owl-s/

The AgentcitiesUK Network

Julian Padget
University of Bath
United Kingdom
agentcitiesuk-request@agentcitiesuk.net

The EPSRC (UK's Engineering and Physical Sciences Research Council) is funding agentcitiesUK as a national network for the next three years (December 2003 to November 2006). Being a *network*, almost all of its budget is for organising events and paying travel and subsistence. As anyone who has coordinated a network will tell you, getting the money is the easy part: the challenge is to get the community engaged.

The network's broad aims are to enable:

- Diffusion of the outputs from the European Union funded Agentcities projects and various projects that have followed on from them, such as Agentcities.ES in Spain, the @lis technology net (EU and Latin America), the ANEMONE project in Italy, and Agent.Enterprise and Agent.Hospital in Germany (see AgentLink News issue 14 of May 2004 for more details on these).
- Development of infrastructure and an enabling framework for a nationally distributed laboratory for agent experiments, which builds on the work that has been done recently on the Agentcities/opennet testbed (again, see AgentLink News issue 14 or http://www.x-opennet.org for more details).
- Provision of a national and a European context for masters and doctoral research in agent-based services through liaison with AgentLink III and other European agent research activities, the Challenge Days (see below) which focus on particular problem domains, and through bursaries for research students to make short visits to other research groups in the UK.

In order to do this, we have to find ways to encourage the community to participate in various kinds of events and activities. The typical vehicle for this is a workshop, but although there will probably be at least one workshop supported by AgentcitiesUK — most likely co-located with a conference in the UK — we thought we would take the risk of trying out a couple of less conventional activities:

1. **Research student mobility**: not new, of course, as other EU projects have done this, but perhaps novel for a UK network. The details and application information are on the network website, but in summary, the idea is to provide a bursary (capped at £500/week to cover travel and subsistence) to support a research student to visit another agent research group — university or industry, indeed the latter is very much encouraged — either because it would be mutually beneficial for the research activities of both, or to facilitate knowledge transfer. The application procedure is intentionally lightweight, just requiring a short case describing the activity and supporting letters from supervisor(s) and host(s). After the visit, a short report will also be needed.

2. **Challenge Days:** these are a real experiment and in nature will probably be rather like the first requirements gathering phase of a software project. Whatever happens, we will be learning as we go along. Each one will focus on a problem domain in which there are significant challenges for conventional software as much as for agent-oriented software; in the first instance we are looking at eHealth. The aim is to look at the selected problem domain, with the help of several "problem owners", from the perspective of agents and agent-oriented software engineering, so as to discover strengths and weaknesses of the agent approach, and above all to discover new solutions to existing problems.

The nature of each Challenge Day will depend very much on the participants and on the topic area, but some of the main goals are:

- To evaluate techniques and tools practically in an expert and open forum, assisting the establishment of best practices and spreading understanding.
- To provide technology vendors with a forum to demonstrate the added value and usefulness of their products.
- To provide motivating examples to promote understanding of the commercial opportunities.
- To bring together industrialists and academics to foster a shared view of the priorities of the field.
- To create an environment from which new research projects might emerge.
- To provide roadmap and technology feedback information to the EPSRC from an event involving key players from the UK multi-agent systems community.
- To energise activities in the open agent deployment field.
- To provide a training and technology transfer forum.

Although as indicated above, the structure of the Challenge Days will be flexible, the rough format will be that the "problem owners" will initiate the process by each describing a selected issue and then groups formed from the participants will go and brainstorm on the problem, produce initial solutions, and then iterate, culminating in a short report on the issues and the potential solutions. The meetings will take place over two days, including time in the evening for more relaxed exploration of the problem. To facilitate the process, there will be permanent internet access and other computing facilities (participants will be expected to bring a laptop, wireless or otherwise). The activity could be likened to the first stages of rapid/extreme development — indeed, some groups may go so far as to produce prototypes

— but of greater value will be the evaluation of the activity and its presentation as a kind of "white paper".

The call for participation for Challenge Day 1, which is on eHealth, appears below. Meanwhile, we are also starting the organisation of Challenge Day 2, for which the topic is eGovernment and eDemocracy. Time and location are not fixed, but we are planning for Spring 2005. If you are interested in helping run that event, please get in touch. Likewise, if you have a topic that you think should be given a Challenge Day treatment, and/or you would like to host one of the Challenge Day meetings, please get in touch.

To give a flavour of some potential topics, here are some suggestions to stimulate reaction (we intend to run between 4 and 5 workshops, depending on budget):

Agentcities.UN: An information community where users can link together, classification of questions, fora for exchange of ideas, managing debates and decision making — translation, summarisation, dissemination.

Agentcities.NHS_Direct: Health and welfare services supported and integrated on an agentcities network.

Agentcities@skool!: Campus management and educational services, virtual universities, academic collaboration, paper-writing assistants, scholarly agents (similar to CiteSeer).

Agentcities-wild: Agentcities attacks, defences, security, metrics, instrumentation, administration.

Agentcities GRID: Grid service setup, interfaces, performance, accessibility, interchange.

Engineering Agentcities: development methodologies, reusable libraries, training courses, supporting technologies.

The most distinctive characteristics of agent-oriented software, namely autonomy, reactivity, proactivity, and sociability are also the very same features that make conventional software engineers and even more so conventional management worried about taking an agent-oriented approach. Another outcome of the Challenge Day discussions, it is hoped, will be arguments about the value of agents and how they can deliver more than conventional software — at least, let's hope that is the case!

After this overview of the AgentcitiesUK network, we look forward to being able to report more when we have held a few activities. Please participate: come to the Challenge Days, and use the mobility funding.

The author is pleased to acknowledge that several members of the UK agent community wrote supporting letters to the EPSRC in respect of the funding application.

AgentcitiesUK.net Challenge Day 1: eHealth

22nd – 23rd September 2004
University of Bath
Bath, UK

The first AgentcitiesUK.net challenge day will be held on the 22nd and 23rd of September 2004 at the University of Bath.

The purpose of a "Challenge Day" is both to extrapolate from current research and to brainstorm on completely new possibilities for the use of agents. The focus is on how agents can help us do tasks that at present cannot be accomplished, rather than automating existing practices. The aims are to identify real agent added-value and develop sustainable arguments that autonomy is an opportunity, not a threat.

The goal of the first AgentcitiesUK challenge day specifically is to take this approach to the application of agent technologies in the UK healthcare sector. It will be attended by researchers and postgraduate students from the agents community, "problem owners", and people who work in both sectors.

To apply for a place and read the details of financial support provided please visit http://agentcitiesuk.net/ and fill out the event registration form.

Late requests to participate can be directed to the organisers.

Latest News from the Standardisation World

Monique Calisti
Whitestein Technologies AG
Switzerland
mca@whitestein.com

As announced in the previous edition of the AgentLink newsletter, the "Latest News from the Standardisation World" column is an initiative, promoted by the AgentLink III project, which aims to report regularly on relevant standardisation activities and events.

Are you directly involved in any standardisation activity? Are you aware of important standardisation work and events that you would like to publicise? You can contribute by sending your material to: standardisation-activity@agentlink.org.

This month, the latest news summary comes from two international standardisation bodies: the Object Management Group (OMG) and W3C Web Services activities.

OMG — Agent Special Interest Group (SIG)
Under recommendation of the OMG's Analysis and Design Task Force and the Agent SIG, the OMG has just issued a "Request for Information (RFI) on Modeling Agent-based Systems." This RFI is a general request for the computer industry, academia, and any other interested parties to submit information regarding considerations for an agent-based UML. For those who are interested, the RFI can be downloaded from:
http://www.omg.org/cgi-bin/doc?ad/2004-08-05.

The primary purpose for the RFI is to prepare for an agent-notation standardisation effort. After analysis of the RFI responses, an RFP (request for proposal) will be issued to elicit proposals for a unified agent-based modelling language. Such a modelling language would be expected to extend UML 2.0 wherever possible. However, this requirement would not be mandatory. The RFP is expected to be issued in Autumn 2005. Any questions about this process or how to get involved should be addressed to OMG's Analysis and Design Task Force and Agent SIG chair, James Odell (omg@jamesodell.com).

Contribution by:
James Odell
email: omg@jamesodell.com
Agentis Software, Inc.

W3C — XMLP Working Group
The W3C XMLP Working Group has reached Last Call status with the SOAP Message Transmission Optimization Mechanism (MTOM) specification [1], which is an extension to the SOAP 1.2 Recommendation [2]. MTOM provides a common means of optimising the transmission and wire format of SOAP messages. MTOM is built on the XML-binary Optimized Packaging (XOP) convention [3], a means of more efficiently serialising XML Infosets.

The W3C Web Services Description Language (WSDL) 2.0 specification [4] is nearing Last Call status (expected in the coming weeks). WSDL 2.0 promises better interoperability and more flexibility than WSDL 1.1, its predecessor. WSDL 2.0 supports SOAP 1.2 and offers interface inheritance and an extensible set of Message Exchange Patterns (MEPs). It is also more Web friendly, as operations can be marked as "Safe" and HTTP is better supported.

The W3C Web Services Choreography Working Group has released the First Public Working Draft of its Choreography Description Langauge (CDL) [5], an XML–based language for defining multi-party, peer-to-peer collaborations of Web Services independently of hosting environment or platform.

The W3C Web Services Activity recently announced a call-for-participation in an upcoming workshop on Constraints and Capabilities for Web Services [6]. This workshop will explore ways to permit Web services to express their service capabilities, constraints on their use, policies, and other metadata in an interoperable, machine-processable form. The workshop is free and open to both W3C members and non-members. However, space is limited and each participant must submit a position paper. Registration details are available at:
http://www.w3.org/2004/06/ws-cc-cfp.html#Expected

Contribution by:
David Booth
email: dbooth@w3.org
W3C Fellow, HP Software

References

1. MTOM: http://www.w3.org/TR/soap12-mtom/
2. SOAP 1.2: http://www.w3.org/TR/soap12/
3. XOP: http://www.w3.org/TR/xop10/
4. WSDL 2.0: http://www.w3.org/TR/wsdl20
5. CDL: http://www.w3.org/TR/ws-cdl-10/
6. Workshop on Constraints and Capabilities for Web Services: http://www.w3.org/2004/06/ws-cc-cfp.html

Communication about Ontological Discrepancies

Robbert-Jan Beun **Rogier M. van Eijk**

Utrecht University
The Netherlands
{rj,rogier}@cs.uu.nl

When two agents exchange information, they need a conceptualisation of the domain of interest and a shared vocabulary to communicate facts with respect to this domain. The two conceptualisations can be expressed in so-called ontologies, which are often defined in a formal language, such as a programming language or a formal logic. An ontology is a catalog of the types of things that are assumed to exist in a domain of interest *D* from the perspective of a person who uses a language *L* for the purpose of talking about *D* [3]. Hence, it abstracts the essence of the domain of interest and helps to classify and distinguish the various types of objects in the domain, their properties, and relationships.

In *collaborative* performance of tasks, agreement between different agents with respect to their ontologies is crucial or, at least, the agents should be aware of existing discrepancies. Various troublesome situations may arise due to ontological discrepancies. One of the agents, for instance, may by mistake assume that a particular property is applicable to particular types of objects, while the other has no knowledge of that. Ontological discrepancies may cause serious communication flaws, and the generation of adequate feedback in order to repair such flaws is an essential part of modelling a proper communication process.

In principle, there is a range of approaches to achieve ontological agreement. In our research [1], we aim at the design of a mechanism that solves ontological discrepancies during the communication process. This recognises the fact that agents may autonomously develop their own ontologies, so discrepancies need to be resolved on-the-fly. In order to design such mechanism, we make three assumptions, as follows:

1. the two ontological representations of the participating agents are independent and may, therefore, be formalised in different languages;
2. the information in a received message is in agreement with the ontology of its sender — in other words, the communication channel is ideal and there are no translation failures from the sender's ontology to its communication vocabulary;
3. ontological representations can be evaluated in a common domain of interest, i.e., both agents are able to observe the common domain they refer to when communicating.

A significant problem is that, given a particular ontological discrepancy, an agent has abundant possibilities for feedback. Suppose, for instance, that agent A asks the question: "Is this file running?", and that B's ontology contains the information that files are a particular type of items and that running is only applicable to processes. Then, examples of possible ways in which B may react are:

Ontological *discrepancies* may cause serious communication flaws, and the generation of adequate feedback in order to repair such flaws is an essential part of modelling a proper communication process.

B1: "Sorry, I don't understand that"
B2: "Running is not applicable to files"
B3: "Running is only applicable to processes"
B4: "Files are not processes"
B5: "Is running applicable to files?"
B6: "I did not know that running is applicable to files"
B7: "What do you mean by running?"

In the communication process, therefore, we clearly distinguish between message interpretation and message generation.

In the interpretation process, background assumptions (called "presuppositions") are extracted from the message and subsequently compared with the receiver's ontology. The receiver may detect discrepancies by, for instance, type conflicts, ontological gaps, and particular inconsistencies that emerge during the conversational process. For various reasons, we have decided to represent the sender's ontology as a type-theoretical context. The decision criteria for discrepancies are thus expressed in terms of entailment in type theory, where the addition of particular information to a given context yields, in the technical jargon, legal or illegal contexts. In the interpretation process, presuppositions from the message are considered as additional information and compared to the type-theoretical context of the receiving agent.

Message generation is still in a rudimentary phase, but we believe that at least the following information should be included for adequate feedback utterances to be generated: (i) the kind of ontological discrepancy, (ii) the role the agents play in the conversation, (iii) the agents' beliefs and their common beliefs about *D*. The message generator produces two types of information: first, the problem itself — i.e., which information causes the discrepancy — and second, a possible way to solve the problem, for instance, a question about the interlocutor's ontology or a statement about its own ontology.

Essential to our approach is that ontological discrepancies are treated at the level of agents themselves, without the aid of an external observer, and therefore on the basis of their own subjective view on the world. In other words, there is no reference to any (implicit) third ontology. It also means that the framework abstracts from a notion of truth which is inherent to model-theoretic approaches. Agents work towards agreement on the basis of their belief states and communicative acts.

Future work will include an extension of the basic framework to richer ontologies and to a richer communication language. Probably more important is the development of a more fundamental view on the generation of adequate feedback, based on, for instance, a formalisation of Gricean cooperation maxims [2] and the information mentioned above, such as roles and various types of beliefs held by the agents.

References:

1. R.J. Beun, R.M. van Eijk, and H. Prust. Ontological Feedback in Multiagent Systems. In C. Sierra and L. Sonenberg, editors, Proceedings of third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), pages 110 – 117, ACM Press, 2004.
2. H. Grice. Logic and conversation. In P. Cole and J. Morgan editors, Speech acts. Syntax and Semantics, Vol. 11, pages 41 – 58. Academic Press, New York, 1975.
3. J.F. Sowa. Knowledge representation: Logical, philosophical, and computational foundations. Brooks/Cole, Pacific Grove, CA, 2000.

Sociologically Inspired Engineering

David Hales
University of Bologna
Italy
dave@davidhales.com

Bruce Edmonds
Manchester Metropolitan University
United Kingdom
bruce@cfpm.org

It is always hard to predict the kinds of cross fertilisations that will occur between different areas within the interdisciplinary hot-house of agent-based research. It is now widely established that biologically-inspired mechanisms are applicable as basis for the engineering of distributed and robust systems. For example, many adaptive agent algorithms draw inspiration from biological evolution, and there has been much recent work on ant-inspired routing methods to name just two. We discuss here recent results of cross fertilisation within agent-based research that is less widely known: recent developments by which mechanisms derived from social theories, particularly those expressed computationally through agent-based social simulation, are being applied to tough engineering problems in distributed settings.

The idea of using "social metaphors" for thinking about self-organising software is not new. Marvin Minsky's classic A.I. text "Society of Mind" [5] explicitly envisages minds as composed of semi-autonomous entities with coalitions, conflicts and hierarchies. What is new is the application of techniques from multi-agent-based social simulation (MABS) to the hard distributed engineering problems that agent researchers are currently attempting to solve.

For well over a decade, simulation-friendly social scientists have been using multi-agent-based simulations to develop, test and communicate mechanisms of social emergence. This means programming individual agents with (often sophisticated) individual behavioural and learning rules and then executing simulations to determine what kinds of societies and structures emerge. The focus of much of this research has been on trying to gain a greater understanding of human societies — the world we actually live in. Interestingly, some of the mechanisms that have been discovered are directly relevant to open engineering issues in distributed systems.

As was described in AgentLink News 14 by Steve Phelps and Peter McBurney [7], ideas from computational economics are informing new automated market design algorithms for deployment in multi-agent systems. Many of these ideas got started when progressive economists, attempting to understand what methods real people use to solve such problems, started making computational models of such methods.

> ...recent developments by which mechanisms derived from social theories, particularly those expressed computationally through agent-based social simulation, are being applied to tough engineering problems in distributed settings.

A general problem that pervades much of agent research is that of maximising system-level performance while allowing individual agents reasonable levels of autonomy. In many situations there arises a contradiction between these two aspects. This kind of dilemma happens in human societies all the time, for example, when someone decides not to pay on a short train ride (free-ride) or evade tax by not declaring income. One way to stop these behaviours is to impose draconian measures via centralised government control — ensuring all individuals behave for the common good — so as to stop free riders. However, this is costly and hard to police, as well as it raises other issues such as who polices the police? In the parlance of agent-oriented systems: the method does not scale well, is sensitive to noise, and has a high computational overhead.

In the context of actually deployed massively distributed software systems, Peer-to-Peer (P2P) file sharing applications (such as KaZaA and eDonkey) have similar problems — most users only download files rather than sharing them. This limits the effectiveness of such systems. Even when the P2P client software is coded to force some level of sharing, users may modify and redistribute a hacked client. It has been noted that P2P file sharing is one of the applications in which only a small number of altruists are needed to support a large number of free riders. Consequently, it can be argued that this might be why popular P2P applications tend to be limited to sharing only files rather than processor or distributed storage, for example.

This sort of case can be seen as an example of a more 'fundamental' issue: how can one maintain cooperative (socially beneficial) interactions within an open system under the assumption of high individual (person, agent, or peer) autonomy. An archetype of this kind of social dilemma has been developed in the form of a minimal game in which two players each selected a move from two alternatives and then the game ends and each player receives a score (or pay-off), the Prisoner's Dilemma game (see below, if you are not familiar with it).

The Prisoner's Dilemma

Table 1 shows a so-called "pay-off matrix" for the two-player game. If both cooperate then both get a reward, the score R. If both defect they are punished, the score P. If one player defects and the other cooperates then the defector gets T (the temptation to defect), the other getting S (the sucker). When these pay-offs, which are numbers representing some kind of desirable utility (for example, money), obey the constraints $T > R > P > S$ and $2R > T+S$, then we say the game represents a Prisoner's Dilemma. When both players cooperate this represents maximising of the collective good but when one player defects and another cooperates this represents a form of free-riding. The defector gaining a higher score (oh, the temptation!) at the expense of the co-operator (who then becomes a sucker, really).

	cooperate	**defect**
cooperate	R, R	S, T
defect	T, S	P, P

Table 1. The Prisoner's Dilemma Game: two players select either to cooperate or defect.

A game theoretic analysis of the prisoner's dilemma drawing on the Nash equilibrium solution concept (as defined by the now famous John Nash) captures the intuition that a utility-maximising player would always defect in such games because, whatever the other player does, a higher score is never attained by choosing to cooperate. The Nash Equilibrium might be a partial explanation for why there is so much free-riding on existing P2P file-sharing systems — users are simply behaving to maximise their utility. However, do we have any way to solve this problem without going back to centralised control or closed systems? A Nash analysis gives us a good explanation for selfish behaviour but not for altruistic behaviour. As stated earlier, even in P2P file sharing systems there are some altruists (keeping the show on the road).

It has been argued by many researchers from the social and life sciences that human (and even animal) societies produce much more cooperation than a Nash analysis would predict. Consequently, various cooperation promoting mechanisms (often using the prisoner's dilemma as their test bed) have been proposed over the last half-century. Within the last decade, there has been a flood of new results in the area due to the increasing availability of computational power and the advent of agent-based social simulation.

This work is characterised more by the experimental method than by the development of a *prior* formal theory. This is because the real distributed systems that are the target of this work are complex and only partially known, forcing theory to be a matter of falsifiable hypothesis and experiment. Thus this work is often closer in method to the natural sciences than to the formal sciences [3].

For well over a decade, simulation-friendly social scientists have been using multi-agent-based simulations to develop, test and communicate mechanisms of social emergence.

We now concentrate on describing two mechanisms, imported from the social simulation literature, that have been directly applied to the problem of maintaining high levels of cooperation in file-sharing P2P applications. One is actually deployed and working; it is called BitTorrent and uses mechanisms popularised by the political scientist Robert Axelrod. The other, offering a potentially more robust and scalable method, applies some very recent ideas from social simulation within a prototype simulation environment.

BitTorrent [2] employs a form of the tit-for-tat strategy popularised in the 1980's by computer simulation tournaments of an iterated version of the prisoner's dilemma. Researchers were asked to submit programs (agents, if you like) that repeatedly played the game against each other [1]. The result of all these tournaments was that a simple strategy called tit-for-tat did remarkably well against the majority of other submitted programs (although other strategies can also survive within the complex ecology that occurs when there is a population of competing strategies).

Tit-for-tat (TFT) operates in environments where the prisoner's dilemma is played repeatedly with the same partners for a number of rounds. The basic strategy is simple: an agent starts by cooperating then, in subsequent rounds, copies the move made in the previous round by its opponent. This means defectors are punished in the future: the strategy relies on future reciprocity. To put it another way, the "shadow" of future interactions motivates cooperative behaviour in the present. In many situations this simple strategy can outperform pure defection.

In the context of BitTorrent, a TFT-like strategy plays a significant role in ensuring cooperation. The basic BitTorrent mechanism is simple: files are split into small chunks (about 1MB each) and downloaded by peers, initially from a single hosting source. Peers then effectively "trade" such chunks with each other using a TFT-like strategy — i.e., if two peers offer each other a required chunk then this equates to mutual cooperation. However, if either does not reciprocate then this is analogous to a defect, and the suckered peer will retaliate in future interactions.

The process is actually a little more subtle because each peer is constantly looking at the upload rate / download rate from each connected peer in time — so it does not work just by file chunk but by time unit within each file chunk. While a file is being downloaded between peers, each peer maintains a rolling average of the download rate from each of the peers it is connected to. It then tries to match it's uploading rate accordingly. If a peer determines that another is not uploading fast enough, then it may "choke" (stop uploading) to that other.

Additionally, peers periodically try connecting to new peers randomly by uploading to them – testing for better rates. This means that if a peer does not upload data to other peers (a kind of defecting strategy), then it is punished by those other peers which will not share file chunks with the "defector" in the future — hence a TFT-like strategy based on punishment in future interactions is used.

Axelrod used the TFT result to justify sociological hypotheses — such as understanding how fraternisation broke out between enemies across the trenches of World War I. Cohen has applied a modified form of TFT to produce a file sharing system resistant to free-riding. However, TFT has certain limitations, it requires future interactions with the same individuals, and each has to keep records of the last move made by each opponent. What about situations in which many interactions are with strangers with no promise of future

interaction and/or take place between millions of nodes? This is often the case in large-scale file sharing systems and might be true for other kinds of P2P applications such as storage and processor sharing. In such circumstances, TFT-type strategies are of limited use — does the result that in a single-round game it often pays to free-ride mean that such systems are doomed to be taken over by free-riders?

Recent work, drawing on agent-based simulations of cooperative group formation based on "tags" (social labels or cues) and dynamic social networks suggests a new mechanism which does not require reciprocal arrangements. It is based on the idea of cultural group selection and the well-known social psychological phenomena that people tend to favour those believed to be similar to themselves — even when this is based on seemingly arbitrary criteria (e.g. supporting the same football team). Despite the rather complex lineage, like TFT, the mechanism is refreshingly simple. Individuals interact in cliques (subsets of the population). Periodically, if they find another individual who is getting higher utility than themselves they copy them — changing to their clique and adopting their strategy. Also, periodically, individuals form new cliques by joining with a randomly selected other.

Defectors can do well initially, suckering the co-operators in their clique — but ultimately all the co-operators leave the clique for pastures new — leaving the defectors all alone with nobody to free-ride on. Those copying a defector (who does well initially) will also copy their strategy, further reducing the free-riding potential in the clique. So a clique containing any free-riders quickly dissolves but those containing only co-operators grow.

Given an open system of autonomous agents, all cliques will eventually be invaded by a free-rider who will exploit the clique hence dissolving it. However, so long as other new cooperative cliques are being created, then cooperation will persist in the overall population. In the context of social labels, or "tags", cliques are defined as those individuals sharing particular labels (e.g., supporting the same football team). In the context of P2P systems, the clique is defined as all the other peers each peer is connected to (i.e., its neighbourhood) and movement between cliques follows a process of network "re-wiring".

Through agent-based simulation, the formation and maintenance of high levels of cooperation in the single round prisoner's dilemma and in a P2P file sharing scenario (containing over 100,000 peers) can be tested. This has been done and these mechanisms have been shown to be effective [4]. The mechanism appears to be highly scalable with zero scaling cost — i.e., it does not take longer to establish cooperation in bigger populations. There are a number of outstanding issues that need to be addressed before this new technique can be deployed. However, as previously mentioned, a deployable system would have applications beyond file sharing.

A general problem that pervades much of agent research is that of maximising system-level performance while allowing individual agents reasonable levels of autonomy.

Here we have outlined just two mechanisms, arising from social science via agent-based social simulation, targeted towards P2P applications. However, there are many other areas in which the same kind of cross-fertilisation is in the process of occurring. These include areas such as trust, reputation, market processes, group formation, normative regulation, negotiation, and even emergent ontologies. For a flavour of recent agent-based social simulation work, take a look at the *Journal of Artificial Societies and Social Simulation (JASSS)* [6], as well as the MABS and ESOA series of workshops (both published in Springer's LNAI series). A workshop currently being organised will look specifically at the use of metaphors and ideas originating in social phenomena to the engineering of complex computational systems (details below).

Increasingly, software engineers are working with distributed and noisy systems through necessity. They are tackling, essentially, sociological and economic questions — though many do not realise this. Agent-based simulation offers a large body of algorithmically specified mechanisms and techniques that can go a long way towards solving these problems. In the AgentLink II roadmap, engineering with social metaphors was placed as a medium to long-term possibility. It seems that this is developing more quickly than expected. Increasingly, we may see cutting-edge software engineering and social science intersect, through mutual self-interest, into a single grouping or, dare we say, clique. You will be able to recognise when this has occurred — a clique composed of sociological computer scientists and computer-literate sociologists is likely to be tagged by being fairly scruffy!

References:

1. Axelrod, R. (1984). The evolution of cooperation. Basic Books, NY.
2. Cohen, B. (2003). Incentives Build Robustness in BitTorrent. Presented at the 1st Workshop on the Economics of Peer-2-Peer Systems, June 5-6, 2003, Berkley, CA. (Available at:www.sims.berkeley.edu/research/conferences/p2pecon/).
3. Edmonds, B. and Bryson, J. (2004). The Insufficiency of Formal Design Methods — the necessity of an experimental approach for the understanding and control of complex MAS. Presented at AAMAS 2004, July, New York. (Available at: http://www.aamas2004.org).
4. Hales, D. (2004) From selfish nodes to cooperative networks - emergent link based incentives in peer-to-peer networks. Proc. of the 4th IEEE International Conference on Peer-to-Peer Computing (P2P2004), IEEE Computer Soc. Press.
5. Minsky, M. (1988). Society of Mind. Simon & Schuster.
6. Journal of Artificial Societies and Social Simulation (JASSS), available at: (http://jasss.soc.surrey.ac.uk/JASSS.html).
7. AgentLink Newsletter 14, research overview article by Steve Phelps and Peter McBurney, available at: (http://www.agentlink.org/newsletter).

Electronic Institutions and Agents

Ulises Cortés
Universitat Politècnica de Catalunya UPC
Spain
ia@lsi.upc.es

The study of agent-mediated electronic institutions (e-Institutions for short) is a new and promising field where the interactions among a group of (software) agents are ruled by a set of explicit or implicit norms. Such norms are expressed in a computational language representation that agents can interpret. They should not be considered as a negative constraining factor but rather as an aid that guides the agents' choices and reduces the complexity of the environment, thus making the behaviour of other agents more predictable.

The approach proposed by Noriega [6] was the first to introduce an abstraction of the economic-theoretic notion of *institutions*. In his work, Noriega coined the term *agent-mediated electronic institution* to describe the construction of computational environments that allow heterogeneous agents to successfully interact by imposing appropriate restrictions on their behaviour. From a social point of view, e-Institutions have been proposed as a formalism to define and analyse agent protocols with a view to achieving global and individual goals [7]. From Noriega's work to the actual lines of research, represented by systems such as HARMON*IA* [8], ISLANDER [3], and OperA [1], or approaches such as those by Fornara [4] and López y López [5], the study of e-Institutions is taking recently a relevant and active role in the research agenda of the European agent community.

However, there are several issues yet to be resolved. One of the main issues is that current work is either too theoretical, focused on norm formalisation by means of very expressive logics that are computationally hard, or too practical, focused in the implementation of e-institutions but losing accuracy and expressiveness in the normative system. No satisfactory connection between both approaches has been defined. Another issue that researchers in the field face is the variety of terminologies that are currently in use. In Figure 1, I try to show a classification of the actual research tendencies in the field: the vertical axis indicates the way norms are represented in each system (adapted from [8]). Finally, an additional issue is that, despite the fact that a significant amount of theoretical and practical work has been carried out in defining e-Institutions, as well as their formalisation and implementation, they have been applied to small or quite simple experimental setups. Very few projects have resulted in real applications, as has MASFIT (http://www.masfit.net).

In the following, I shall review briefly five Ph.D. theses, all of which are related to Agents and Institutions. The theses were presented during 2003 in different European universities, showing the strength and activity of this area of research. The respective systems are presented in alphabetical order.

Agents & Institutions

HARMON*IA* introduces a new framework for electronic organisations that defines a multilevel structure, from the most abstract level of the normative system to the final multi-agent implementation. The thesis, which received the 2003 ECCAI Artificial Intelligence Dissertation Award, discusses the main issues surrounding the implementation of *norms* in agent-mediated institutions. The main observation is that norms are specified in regulations that are described (on purpose) at a high abstraction level. The level of abstraction is high so as not to be dependent on a circumstantial implementation of the norm. Norms should be stable for many situations and for a relatively long time. Therefore, it is obvious that norms do not contain the concrete means for their implementation. To allow for the implementation of norms, HARMON*IA* has a formalism in which one can explicitly specify *how* the norms are translated into concrete norms applied in the context of a concrete organisation. Using these concrete norms, one can then translate them in operational representations, such as rules or procedures, thus indicating *how* norms are to be implemented in the *e*-organisation [8].

ISLANDER is both a formalisation of an e-Institution and a powerful tool that helps engineers on the process of specifying, verifying, and developing real-world multi-agent system applications [3]. It has received several awards, among them the *best prototype paper* at AAMAS02. Esteva argues that open multi-agent systems can effectively be designed and implemented as agent-mediated e-Institutions where heterogeneous agents can participate, playing various roles and interacting by means of the regulated exchange of illocutions; this clearly follows and extend the work done in [6]. In this approach, agent activities within an e-Institution are structures in *conversations* which determine the valid interactions between agents and represent the context where exchanged illocutions must be interpreted. The validity of a conversation is determined by the e-Institution.

ISLANDER follows and culminates the work initiated in [7]. An important feature of ISLANDER is that specifications of *e*-Institutions are programming-language independent.

The OperA Model develops a framework for agent societies to legitimate the concept of autonomy between society requirements and agent goals. Dignum assumes that organisational engineering of multi-agent systems cannot assume that participating agents *will act* according to the needs and expectations of the system design.

The development of OperA includes the formalisation — with formal semantics based on temporal deontic logic — and the design of a methodology where the interaction between agents is represented in such way that: (i) it is independent of the internal design of the agents; (ii) it distinguishes organisational characteristics from agents' own goals (this aim also appears

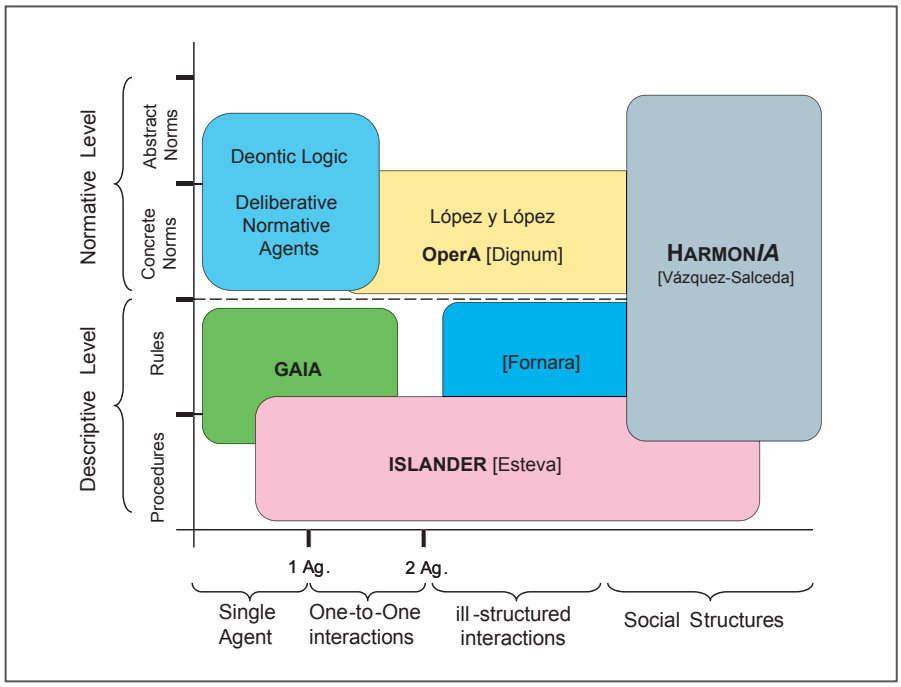

Figure 1 Classification of Research in Electronic Institutions and Agents

in [5]); (iii) it creates dynamic links between organisational design and agent populations; and (iv) it allows for the adaptation of interaction patterns to the characteristics of specific populations [1].

Fornara [4] proposes an operational specification for the definition of a standard agent communication language based on the notion of social commitment. To some extent, this work addresses the issue of *illocutions*, which are essential for the coordination of actions as discussed in [6]. One of the main contributions of this work is the operational definition of the semantics of a verifiable and objective agent communication language based on the notion of social commitment. In this proposal, the semantics is *public*, in the sense that any third party agent that observes the message flow has to be able to draw similar inferences from the interaction, and *objective* in that everyone confers the same meaning to the exchanged messages. This semantic is external with respect to the agent's internal structure, flexible and extensible to let agents cope with varying situations, simple to use by agent designers, yet expressive enough. Finally, with respect to the agents' autonomy, they should only have social constraints on their behaviour.

López y López [5] proposes a general model to unify different normative frameworks for agent-based systems. The model includes all the elements needed by autonomous agents to make decisions concerning their *normative behaviour*. This approach has its roots in the SMART agent framework [2].

The model facilitates the understanding and the computational implementation of norms and multi-agent systems regulated by norms and attempts to provide a model for agents to reason about norms. For this purpose, the identification of the *power* an agent can exert is important as it is a means to influence the action of other autonomous agents. In this regard, it is interesting to study how, when, an by whom agents can be constrained under such power, and how to identify such situations. This implies that *power* is dynamic and that agents should reason about this fact so as to determine whether this affects their plans or not. The thesis also includes a classification of dynamic power relationships.

Concluding remarks

In all these theses there is an attempt to formalise the development of *e*-Institutions for real-world applications, each of them present appealing distributed applications in complex and real domains such as *e*-commerce, *e*-business, or *e*-healthcare. Also, all of them start from a formal specification step that allows identifying and formalising their components. An important point of differentiation is the level of abstraction used to start the specification of norms used in institutions (see Figure 1) and also the abstraction level of the final product. As *e*-Institutions are hard to be fully-modelled in detail, few approaches cover the whole problem, from the most abstract norms to rules and their final implementation, and from the individual agent to the agent society.

All these efforts make clear that norms do not contain themselves the concrete means for their implementation or interpretation. Therefore, in order for norms to be implemented, it is necessary to develop a formalism in which one can explicitly specify *how* the (abstract) norms are translated into concrete norms applied in the context of a concrete organisation. Using these concrete norms one can then translate them into operational representations, such as rules or procedures, to indicate how norms are to be implemented in the *e*-organisation. Tools and mechanisms ought to ensure that certain properties of *e*-Institutions hold before they can be enacted (*i.e.*, agents interact following the specified order and types of messages as expected in that *e*-Institution).

Electronic institutions and multi-agent systems technologies are at hand to produce innovations to face the new technological challenges that emerge from open, dynamic, and distributed systems where autonomous and heterogeneous agents, usually having different owners and interests, may interact exchanging services, resources, information, and may form dynamic coalitions (for example, to expand organisational boundaries). What form these innovations will finally take is difficult to assess, but it is clear that we will see a number of real *e*-organisations where heterogeneous agents will interact, and also e-organisations represented by agents interacting under sets of norms.

References

1. V. Dignum, *A model for organizational interaction. Based on Agents, Founded in Logic*, Ph.D. Thesis, Utrecht University, 2003.
2. M. d'Inverno & M. Luck, *Understanding Agent Systems*. Springer-Verlag, 2001.
3. M. Esteva, *Electronic institutions: from specification to development*, Ph.D. Thesis, Universitat Politècnica de Catalunya, 2003.
4. N. Fornara, *Interaction and communication among autonomous agents in multiagent systems*, Ph.D. Thesis, University of Lugano, 2003.
5. F. López y López, *Social power and norms. Impact on Agent Behaviour*, Ph.D. Thesis, University of Southampton, 2003.
6. P. Noriega, *Agent Mediated Auctions. The Fishmarket Metaphor*, Ph.D.Thesis, Universitat Autònoma de Barcelona, 1997.
7. J. A. Rodríguez-Aguilar, *On the design and construction of Agent-mediated Electronic Institutions*, Ph.D. thesis, Universitat Autònoma de Barcelona, 2001.
8. J. Vázquez-Salceda, *The role of norms and electronic institutions in multi-agent systems applied to complex domains the* HARMON*IA framework*, Ph.D. Thesis, Universitat Politècnica de Catalunya, 2003.

Semantic-Based Knowledge Systems Action Line

Report on Network Partner Projects: Part I

In January 2004, several projects were funded under the "Semantic-Based Knowledge Systems" Action Line, as part of the European Union's Sixth Framework Program (FP6). Whilst many of these projects have a clear focus on Semantic Web research, Agent-based technology plays a key role as an enabling technology for the Semantic Web, as well as an important consumer of much of the content that will be annotated and deployed on the Semantic Web. To facilitate and increase shared awareness, AgentLink has agreed to co-operate with five other Networks of Excellence under this Action Line so as to promote high quality, agent-based technology within the Semantic Web community, and to raise greater understanding within the Agent community of one of the more significant technologies to emerge in the new millennia.

Each issue of the AgentLink newsletter will feature informative articles pertaining to some aspect of research from these partner networks, as well as providing information about upcoming, synergistic activities supported by the Networks in this Action Line (including AgentLink). In this, the first of a two-part feature on the Semantic-Based Knowledge Systems Action Line, three of the five partners are introduced:

- *REWERSE* — Reasoning on the Web
- *MUSCLE* — Multimedia Understanding through Semantics, Computation and Learning
- *KnowledgeWeb* — Semantic Web Enabled E-work and E-commerce

As well as featured articles, links and resources will be shared between the partners (such as the recently published AIM@Shape newsletter), and a shared calendar of events will evolve, to encourage further collaboration and synergy between the Agent and Semantic Web communities. More information on this collaboration, as well as shared resources, can be found at : www.agentlink.org/resources/internetwork-activities.html.

EU FP6 Semantic-Based Knowledge Systems Action Line Partner Projects

KnowledgeWeb
Contact: *Valentina Tamma*
URL: http://kw.dia.fi.upm.es/semanticportal/jsp/frames.jsp

MUSCLE
Contact: *Remi Ronchaud*
URL: http://www.muscle-noe.org/

AIM@SHAPE
Contact: *Manolis Vavalis*
URL: http://www.aimatshape.net

KnowledgeBoard 2.0
Contact: Silverio Petruzzellis
URL: http://www.knowledgeboard.com

REWERSE
Contact: *Uta Schwertel*
URL: http://rewerse.net/

AgentLink III
Contact: *Terry Payne*
URL: http://www.agentlink.org

REWERSE

Reasoning on the Web

François Bry and **Uta Schwertel**
University of Munich
Germany
{francois.bry,uta.schwertel}@ifi.lmu.de

> "For the semantic web to function, computers must have access to [...] sets of inference rules that they can use to conduct automated reasoning."
>
> Tim Berners-Lee, James Hendler, and Ora Lassila. The Semantic Web, Scientific American, May 2001.

About REWERSE. REWERSE is a research Network of Excellence on "Reasoning on the Web" which is funded by the European Commission and Switzerland within the Sixth Framework Programme (FP6) under contract number 506779. REWERSE is funded as part of the IST strategic objective "Semantic-based knowledge systems". REWERSE started on 1st March 2004 and will run for four years. The EC and Switzerland will support REWERSE with more than 5 million Euro.

REWERSE involves 27 European research and industry organisations from 14 European countries and about 100 computer science researchers and professionals playing key roles in applied reasoning. The project administration is undertaken by the University of Munich, Germany, led by the project co-ordinator François Bry (Munich), and deputy co-ordinators Hans Jürgen Ohlbach (Munich) and Jan Małuszyski (Linköping). REWERSE's activities and research results can be followed at http://rewerse.net.

Scientific Context and Focus. The objective of REWERSE is to establish Europe as a leader in the area of reasoning languages for advanced Web systems and applications, often referred to as *Semantic Web* systems and applications. The term "Semantic Web" refers to one of the major current endeavours worldwide in Information Technologies. Its goal may be briefly described as enriching the existing Web with meta-data and data processing (and meta-data processing) so as to provide Web-based systems with advanced (so-called intelligent) capabilities, in particular with *context awareness* and *decision support* strengthening a person-centred, everyday use of the Web.

Web circles such as the W3C have become increasingly conscious of the need for functionality- and application-independent reasoning languages as generic building stones of Web and Semantic Web systems and applications.

For example, an advanced Web service could autonomously plan, for a particular person, a weekend trip that takes into consideration previous travel destinations of that person, preferred accommodations, and adapt the schedule of the trip to the work schedule of that person. Offering such a service requires complex reasoning capabilities where knowledge that has not been explicitly mentioned has to be automatically deduced from existing knowledge. Since in most Semantic Web application scenarios the required advanced capabilities call primarily for reasoning (also referred to as logic, deductive, or rule-based), reasoning languages are essential to such Web applications.

Reasoning languages for the Web, although already considered and at times developed as prototypes in a restricted manner in some research contexts, are still not currently available as a technology. Existing Semantic Web languages and/

or reasoning systems such as DAML+OIL and OWL, BPEL4WS, BPML, DAML-S, ConsVISor, JTP, and Triple are developed mostly from functionality-centred (e.g., ontology reasoning or access validation) or application-centred (e.g., Web service retrieval and composition) perspectives. Complementing these activities, REWERSE promotes a *perspective centred on the reasoning techniques* (e.g., forward or backward chaining, tableaux-like methods, constraint-based reasoning, etc.). By adopting this perspective in the REWERSE project, we can recognise the different forms of reasoning needed by various Web systems and applications. This makes it possible for us to devise a *minimal* collection of complementary and interoperable reasoning languages for the Web. Recently, Web circles such as the W3C have become increasingly conscious of the need for functionality- and application-independent reasoning languages as generic building stones of Web and Semantic Web systems and applications. REWERSE aims at fulfilling this need, promoting applied research on this issue of considerable economical importance. Therefore, technologies developed by REWERSE will represent an essential breakthrough for the current Web and the Semantic Web.

Goals. To achieve its objectives, REWERSE aims at structuring and networking a scientific community on reasoning languages on the Web. A major objective of REWERSE is to provide tangible technological bases (that are currently not available) for the industrial development of advanced Web systems and applications. Striving for tangible outcomes, REWERSE aims to:

- develop a coherent and complete, yet minimal, collection of interoperable reasoning languages for advanced Web systems and applications;
- test these languages on context-adaptive Web systems and Web-based decision-support systems selected as testbeds for proof of concept purposes;
- bring the proposed languages to the level of open pre-standards, i.e., suitable for submission to standardisation bodies such as the W3C.

Research. The research activities of REWERSE are divided into several working groups:

- **Rule markup languages** aiming at unified markup languages and tools for reasoning on the Web (Co-ordination: G. Wagner, Eindhoven; G. Antoniou, Heraklion)
- **Policy specification, composition, and conformance** aiming at user-friendly high-level specifications for complex Web systems (Co-ordination: P. Bonatti, Naples; N. E. Fuchs, Zurich)
- **Composition and typing** aiming at methods and rules for Software interoperability in Web contexts (Co-ordination: U. Aßmann, Dresden; W. Drabent, Warsaw)
- **Reasoning-aware querying** aiming at a query and transformation language for the Web with reasoning capabilities (Co-ordination: F. Bry, Munich; M. Marchiori, Venice)
- **Evolution and reactivity** aiming at specifying the evolution of Web-based data repositories (Co¬ordination: J. J. Alves Alferes, Lisbon; W. May, Göttingen)
- **Web-based decision support for event, temporal, and geographical data** aiming at enhancing event, temporal, and location reasoning on the Web (Co-ordination: H. J. Ohlbach, Munich; S. Abdennadher, Cairo)
- **Towards a Bioinformatics Semantic Web** aiming at adding semantics to the Bioinformatics Web (Co-ordination: M. Schroeder, Dresden; R. Backofen, Jena)
- **Personalised information systems** aiming at user-adapted Web information and teaching systems (Co-ordination: N. Henze, Hannover; A. Martelli, Torino)

Activities. In addition, REWERSE promotes activities to disseminate excellence within its research field throughout Europe, including:

- "University education and training" (Co-ordination: J. Małuszyski, Linköping; N. Eisinger, Munich)
- "Technology transfer and awareness" (Co-ordination: T. Geisler; Munich, S. Carro-Martínez, Madrid)

Current status. REWERSE started on 1st of March 2004 with a very successful kick-off meeting in Munich where all working groups presented their research and established co-operations with other working groups. Such co-operation has been further intensified in several face-to-face meetings.

Evidence of both the relevance and the accuracy of REWERSE's objectives for current research can be recognised from the fact that at month 4 of the project already around 30 articles stemming from REWERSE work (partly initiated in the project preparation phase) have been accepted for publication in peer-reviewed venues (cf. http://rewerse.net/publications.html).

By the end of August, REWERSE's working groups will submit the first round of Deliverables. The Deliverables will be made publicly accessible on the REWERSE Website.

Impact. Reasoning languages for the Web are an emerging technology that is currently not available. This technology will soon represent an essential breakthrough for Web systems and applications. Thus, REWERSE promotes research on an issue of a considerable economical importance. By doing so, REWERSE contributes to the international competitiveness of the European industry in an essential field of today's Information Technologies.

REWERSE aims at establishing itself as one of the world leading virtual research centres on reasoning languages and methods for the Web. REWERSE will ensure that this novel technology is fully exploited and translated into real competitive advantages for the European industry.

MUSCLE

Multimedia Understanding through Semantics, Computation and Learning

Eric Pauwels
CWI
The Netherlands
Eric.Pauwels@cwi.nl

Remi Ronchaud
ERCIM EEIG
France
remi.ronchaud@ercim.org

MUSCLE is an IST Network of Excellence, funded under the European Commission's Sixth Framework Programme (FP6), that aims to integrate European scientific excellence in *Multimedia Understanding through Semantics, Computation and Learning*. It bundles the expertise of 42 research groups and is scheduled to run for four years, starting on March 1, 2004. ERCIM EEIG (France) is responsible for the Network's general management, the scientific activity being coordinated by a Steering Committee, headed by Eric Pauwels (CWI, The Netherlands).

The MUSCLE initiative was born out of the growing conviction that there is an urgent need to improve content-based access to multimedia databases in order to realise the overarching vision of a knowledge-driven society. As databases are accruing ever larger amounts of complex mulimedia documents, and networks allow fast

and almost ubiquitous access to an abundance of resources, we find that progress is hampered by the sheer amount and diversity of the available data. It is becoming increasingly clear that access can only be efficient if based directly on content and semantics, the extraction and indexing of which is only feasible if achieved automatically.

...there is an urgent need to improve content-based access to multimedia databases in order to realise the overarching vision of a knowledge-driven society.

To address these issues, the MUSCLE pan-European network of excellence is open to all, including non-EU participants, who can contribute towards this endeavour. The main purpose of this network of excellence is to foster close collaboration between research groups in *multimedia datamining* on the one hand, *and machine learning* on the other, in order to make scientific progress towards the following goals:

- **Moving from modelling to learning**: Harnessing the full potential of machine learning and cross-modal interaction for the (semi-) automatic generation of robust meta-data with high semantic value for multimedia documents. In particular, MUSCLE researchers will develop software tools and research strategies that enable users to move away from labour-intensive, case-by-case modelling of individual applications, and allow them to take full advantage of generic, adaptive, and self-learning solutions that need minimal supervision.
- **Improving interoperability through understanding:** Improving interoperability and exchangeability of heterogeneous and distributed (meta)data by enabling data descriptions at high semantic levels (e.g., ontologies, XML schemata) and by adding inference schemes that can reason about them at the appropriate levels. To this end, MUSCLE researchers will contribute to relevant international standards and protocols.
- **Creation of expressive and adaptive interfaces:** In the same vein, improve human-machine interface by exploring how machine learning can invigorate the creation of expressive, context-aware, and human-centred interfaces that will be able to effectively assist users in the exploration of complex and rich multimedia databases. With regard to these topics, MUSCLE research will contribute to viability studies and proof-of-principle demonstrations.
- **Ensure durable integration and collaboration through the creation of a virtual lab:** Through innovative use of modern multimedia infrastructure, allow for easy and immediate access to people, data and ideas. By doing this, we can complement actual mobility of researchers with engaging virtual presence and interaction, which is essential for establishing effective networking as well as durable integration at a truly European scale.
- **Spreading of expertise and excellence:** Through dissemination, training (e.g. fellowships), and industrial liaison, contribute towards the distribution and uptake of technology by relevant end-users (industry, education, services, etc.).
- **Societal impact**: By accomplishing the above, facilitate the broad and democratic (i.e., obviating the need for special skills) access to information and knowledge for European citizens.

To carry out its programme, MUSCLE relies on the extensive expertise of its many partners, and is managed by ERCIM and CWI.

For additional information on MUSCLE, its members, and its activities, please visit:
http://www.muscle-noe.org/

Knowledge Web

Realizing the Semantic Web



Ying Ding
University of Innsbruck
Austria
ying.ding@uibk.ac.at

Knowledge Web is a FP6 Network of Excellence that aims to support the transition of ontology technology from academia to industry. The network started on 1st of January 2004, with a budget of around 7 million Euro and 18 participants including leading partners in Semantic Web, Multimedia, Human Language Technology, Workflow, and Agents.

The mission of Knowledge Web is to strengthen the European industry and service providers in one of the most important areas of current computer technology: Semantic Web enabled e-work and e-commerce. We will centre our efforts around the outreach of this technology to industry. Naturally, this includes education and research efforts to ensure the durability of the impact and support to industry. Therefore, the main objectives of Knowledge Web are:

- **Outreach to Industry:** Knowledge Web will jointly set up an ontology language, tools, and methods that cover all the major tasks in working with ontologies. Developing standards helps to solve the interoperability problem. It includes benchmarking, compliance testing, usage scenarios, cookbook style textbooks with best practices, and definition of tool environments based on loosely coupled Web Services. We will also push the creation of an Ontology Registration Authority (ORA).
- **Outreach to Education:** The goal will be to establish a Virtual Institute for Semantic Web Education (VISWE), which will act as the principal focus for educational activities on the Semantic Web. In working towards this end, we will build on the experience of several leading European university groups in designing and delivering courses in this area. We expect to provide up to date learning material, curricula, and, ultimately, new degree programs. At the same time, we will enhance the delivery of course materials by making use of novel Semantic Web technologies in combination with more traditional e-learning environments.
- **Coordination of Research:** Knowledge Web will join research efforts on combining Semantic Web with Web Service technology. A strong cooperation with other relevant fields such as multimedia, agent technologies, and human language technologies will be established. The technology developed in those fields will play an essential role in the realisation of the Semantic Web. On the other hand, the Semantic Web itself poses new challenges and influences research developments in those fields. Knowledge Web will establish a Virtual Research Centre to further the coordination of all related research being done in those different areas.

The cooperation with other projects, networks and initiatives will also play a key role in the activities and results of Knowledge Web. To achieve the goals mentioned above, the main activities

planned for the Knowledge Web network, in those three key directions, include:

Industry

- **Assessment of industrial application needs:** further our understanding of the needs, difficulties, and problems of transferring Semantic Web technologies to current industrial systems.
- **Evaluation for technology selection:** survey and profile existing ontologies, tools, infrastructures, and methodologies in the light of thoroughly understood industrial requirements for the adoption of Semantic Web technologies.
- **Technology recommendations:** produce recommendations, guidelines, and standards to help industry organise, design, and implement the move towards Semantic Web technology enabled IT systems.
- **Promotion of ontology technologies:** promote awareness of the added value of the Semantic Web technologies in the increasingly competitive knowledge economy.
- **Cross-network cooperation:** organise cross-disciplinary research in related areas to provide joint educational and promotional efforts, both in academia and industry.
- **Semantic portal infrastructure:** underpin the integration of the activities of the Knowledge Web partners.

Education

- **Foundations for the Virtual Institute for Semantic Web Education (VISWE):** lay the foundations of VISWE by preparing all prerequisites for its formal establishment.
- **Educational content and event provision:** preparing and delivering course material, educational events, summer schools, complete courses, and training programmes.
- **Semantic delivery platform:** develop a Semantic Platform for the delivery and interaction with learning material about the Semantic Web.

Research

- **Scalability:** address the issue of how knowledge processing and ontology-based tools and tool suites can scale to the Web (or an important part of it) to deal with large volumes of information.
- **Heterogeneity:** how to deal with resources using different languages, terminologies, and models.
- **Dynamics:** how to deal with knowledge evolution, knowledge evaluation, and how to reach dynamic consensus on the corresponding ontologies.
- **Semantic Web language extensions:** identify precise requirements for new languages and language extensions, and to represent the interests of the Knowledge Web network in the development and standardisation of such languages and language extensions, including rule-based languages and query languages.
- **Semantic Web Services:** these will provide input and problems to the other (vertical) activities: scalability, heterogeneity, dynamics, and languages. We will investigate a scalable infrastructure for Web-service discovery and orchestration that provides a new infrastructure for e-work, e-government, and e-commerce.
- **Towards a Virtual Research Centre:** which will monitor and organise the research carried out within Knowledge Web, including dissemination activities such as focused workshops, managing a programme of exchange for researchers, and organising the cooperation within Europe as well as with other international initiatives in America, Asia, and Australia.

Network Partner Calendar Sep – Dec 2004

Each AgentLink issue will feature a Network-Partner Calendar listing all forthcoming events associated with each of the network partners. These include internal meetings, Network-Partner organised events, and sponsored events.
A full listing of this calendar can be found at **www.agentlink.org/resources/internetwork-calendar.html**.

KnowledgeWeb	Sep 1-2	Joint event with FOAF: workshop on Friend of a Friend, Social Networking and the Semantic Web	Galway, Ireland
KnowledgeWeb	Sep 2-4	The 11th International Conference in Artificial Intelligence: Methodology, systems, applications and semantic web challenges	Varna, Bulgaria
KB2.0	Sep 5-12	The KM Summer Camp 2004	Oeiras, Portugal
KnowledgeWeb	Sep 9-10	KnowledgeWeb Industrial Cluster Meeting	Trento, Italy
REWERSE & CoLogNet	Sep 9-10	2nd Workshop on Principles and Practice of Semantic Web Reasoning (PPSWR 2004) at The 20th International Conference on Logic Programming (ICLP)	St Malo, France
MUSCLE	Sep 24-25	IBAI Tutorial Days Data-Mining	Leipzig, Germany
AgentLink	Sep 27-29	8th International Workshop on Cooperative Information Agents	Erfurt, Germany
KnowledgeWeb	Sep 27-29	3rd KWeb Research Area Federated Meeting	Manchester, UK
AgentLink	Sep 29-30	CLIMA V Computational logic in multi-agent systems	Lisbon, Portugal
AgentLink	Oct 1	AgentLink Agent Technology Conference	Zurich, Switzerland
AgentLink	Oct 20-22	5th International Workshop on Engineering Societies in the Agents World	Toulouse, France
KnowledgeWeb	Oct 25-29	KW industrial workshop: OTM2004 Industry Program Workshop on business applications of Semantic Web and Grid Computing	Larnaca, Cyprus
MUSCLE	Oct 25-26	International workshop on multidisciplinary image, video, and audio retrieval and mining	Sherbrooke, Canada
MUSCLE	Nov 4-5	MUSCLE Scientific Meeting	Malaga, Spain
KnowledgeWeb & AgentLink	Nov 7	Semantic Web Technology For Mobile And Ubiquitous Applications	Hiroshima, Japan
KnowledgeWeb	Nov 7	Semantic Web Services Workshop, ISWC 2004	Hiroshima, Japan
KnowledgeWeb	Nov 8-10	KM Europe 2004	Amsterdam, The Netherlands
KnowledgeWeb	Nov 9-11	3rd International Semantic Web Conference	Hiroshima, Japan
KnowledgeWeb, MUSCLE	Nov 25-26	Joint event between KW and aceMedia: European Workshop on the Integration of Knowledge Semantic and Digital Media Technologies. Royal Statistical Society	London, UK
AgentLink	Dec 16-17	European Workshop on Multi-Agent Systems (EUMAS04)	Barcelona, Spain

Agent Based Simulation

Helder Coelho
Universidade Nova de Lisboa
Portugal
hcoelho@di.fc.ul.pt

ABS2004 was the fifth edition of a line of scientific events in the domain of multi-agent systems aimed at discussing methodologies, techniques, and applications conjointly.

In the last 10 years, the modelling and simulation processes of industrial and eco-social systems has benefited from recent developments on Multi-Agent Systems in the field of Distributed Artificial Intelligence. Research in the domain of multi-agent system modelling is now significant due to the adequacy of such systems for the design and simulation of complex systems — i.e., systems composed of many entities in non-linear interactions between themselves and with an external environment. See "Mathematical Models vs Agents" (opposite) for a discussion on the class of systems for which agent-based models are more appealing than traditional methods.

The main goal of this year's workshop was to offer a forum for people interested in agent-based modelling and simulation, pointing out the current research on applications and tools for modelling and simulation based on multi-agent systems. The Workshop promoted contacts between practitioners of multi-agent systems on socio-ecosystems for the exchange of experiences on large-scale applications, modelling methodologies, and management of multi-dimensional (both temporal and spatial) as well as multiple-viewpoint models.

Mathematical Models vs Agents

When are Traditional Mathematical Models not Adequate?

Multi-agent systems are used in many domains in which classical mathematical models are not adequate. This is normally the case when:

- the dynamics of the systems are far from the notion of equilibrium;
- the systems are open (i.e., they require the creation and deletion of entities during the course of simulation);
- emergent phenomena have to be explored;
- the entities are heterogeneous (as, for example, in socio-ecosystems, which combine ecological and social dynamics);
- or simply because they provide a more intuitive understanding of the systems than traditional approaches.

The Agent-Based Simulation international workshop series, sponsored by the SCS (the Society for Modeling and Simulation International), was initiated in 2000 and organised for three years in Passau, Germany by initiative of Bernd Schmidt and Christoph Urban (Passau University). In 2003, the 4th ABS edition was organised by Jean-Pierre Muller (CIRAD) in Montpellier, France, and this year's edition took place in Lisbon, Portugal, with the organisational support of the FCUL (Faculdade de Ciências, Universidade de Lisboa).

The 2004 ABS edition showed the vitality of the multi-agent system modelling and simulation community by the diversity of domains to which this approach is currently being applied: social sciences and sociology, traffic and transportation, economy and management, and ecological and environmental sciences.

Altogether, 42 papers were received, of which 25 were accepted as full papers, and 5 as posters. The scientific programme was structured around 7 tracks on: modelling, social sciences, economy, tools, communications, applications, and sociology. Stimulating discussions followed the paper presentations, thus allowing a clear view of what is being done in this area in Europe, the U.S.A., and Brazil.

Knowledge and Games Workshop

Sieuwert van Otterloo
University of Liverpool
United Kingdom
S.van.Otterloo@csc.liv.ac.uk

The first Knowledge and Games Workshop took place in the weekend of 10th and 11th of July, directly after AgentLink's European Agent Systems Summer School. The workshop attracted 29 researchers from all over Europe and beyond. It was a truly interdisciplinary workshop in which logic, computer science, philosophy, mathematics, and economic theory all converged in a series of invited talks, accepted papers, and discussion sessions. The main topic of this workshop was the interaction and overlap between studies of knowledge and game theory, from the point of view of those various disciplines.

Ron van der Meyden (UNSW, Australia) opened the workshop with an invited talk, explaining how the MCK model checker he is developing can be used for verification of several multi-agent protocols. Then Thomas Agotness, from the University of Bergen, Norway, discussed the multi-agent strategic logic ATEL. More precisely, he critically reviewed the first version of this logic and showed that certain claims about this logic did not hold. Alexandru Baltag (University of Oxford) then confronted the audience with an increasingly complex series of diagrams that illustrated how one can truthfully learn that one is being deceived (but for all I know, we might have been deceived ourselves). This logical talk was followed by something completely different. Jeroen Donkers, from the IKAT in Maastricht, The Netherlands, discussed the use of opponent modelling in combinatorial game theory. In his work, he tries to develop a symmetric method of modelling the opponent's strategy: an approach in which both players have a model of each other.

After lunch, Francien Dechesne from the University of Tilburg and the Eindhoven University of Technology, discussed information friendly logic. She looked at the question of whether the Thompson transformations, well-known in game theory, could be applied to the IF-logic of imperfect information games. In the resulting discussion she and Wilfrid Hodges (Queen Mary, University of London) concluded that perhaps this logic was not about games after all. This talk was followed by Sieuwert van Otterloo's introduction of knowledge condition games. Such games, in which agents interact in a protocol so that certain knowledge situations arise, could have been a hugely successful model for multi-agent interaction, had Sieuwert not decided to prove that solving such games is actually NP-hard. The day was then concluded by an invited talk by Hans van Ditmarsch, who looked at the game theory behind the Pit trading game. His analysis showed how this game can be modelled in logic, and showed that the three-player game resembles the classic "chicken" game situation. This prompted the question from the audience as to what extent logic and game-theoretic analyses are related.

On Sunday, the workshop opened with an invited talk by Marc Pauly (IRIT, France), who explained that "a strong coffee or a light talk" were the best ways to start the day. In the ensuing talk (the strong coffee came later), he posed the questions that trouble his conscience when he lies awake late at night: how much logic actu-

ally contributes to game theory, how realistic the sophisticated reasoning common in game theory actually is, and whether the assumptions of epistemic logic are not too simple after all. Thomas Forster (University of Cambridge) then discussed a different type of games: predator-prey games in which agents move continuosly. The result he showed was that such games can be approximated by combinatorial games. Eric Pacuit, from the City University of New York, then moved the topic back into solid logical terrain with his talk on a logic for knowledge and communication. After that, Merlijn Sevenster from the University of Amsterdam took over with a virtual presentation. Due to a knee injury, he was not able to attend the workshop, so he prepared a video presentation on complexity issues in information friendly logic. The main idea was that imperfect information increases complexity, but in his talk it became clear that there are exceptions to this general idea.

After lunch, a paper presentation was given by Wouter Teepe from the University of Groningen, The Netherlands. Teepe offered the audience a choice from two protocols that could be used for determining whether two agents had the same knowledge, without actually revealing what the agents knew. This turns out to have many practical applications, for instance in police investigations. The day was then concluded by Paul Harrenstein's (University of Utrecht, The Netherlands) exposition of a game-theoretical notion of consequence that showed how game theory can be used to define more interesting interpretations of logical formulæ.

There was a general consensus that this workshop, sponsored by Agentlink III and the Department of Computer Science of the University of Liverpool, was very successful. Many participants are hoping that a second workshop will be organised in the near future, perhaps as a yearly event.

Workshop on Deontic Logic in Computer Science

Alessio Lomuscio
University College London
United Kingdom
A.Lomuscio@ucl.ac.uk

DEON04, the Seventh International Workshop on Deontic Logic in Computer Science, was held at the Savoy Hotel, Madeira, Portugal, from the 26th to the 28th of May 2004. The workshop aimed at bringing together researchers interested in topics related to the use of deontic logic in computer science.

The DEON workshop series traditionally promote research on the relationship between normative concepts and computer science, artificial intelligence, organisation theory, and law. In addition to this, DEON04 strived to have a special emphasis on the relationship between deontic logic and multi-agent systems. There have been six previous DEON workshops: Amsterdam, December 1991; Oslo, January 1994; Sesimbra, January 1996; Bologna, January 1998; Toulouse, January 2000; and London, April 2002. Selected papers from each of these workshops have been published in international forums.

The workshop included 18 paper presentations and two invited talks. The topics discussed span from theoretical investigations of deontic concepts and their formalisation in logic to the use of deontic formalisms to verify and reason about multi-agent system applications. We believe this made DEON04 a well-balanced and interesting workshop.

The workshop initiated with an invited talk by Michael Wooldridge (Liverpool) on the use of the ATL logic for modelling social laws and related concepts. Following this, Jan Broersen (Utrecht) presented an article on the use of an extension of a BDI logic in which obligations are expressed as directed modalities for the specification of organisations. This was followed by a presentation by Davide Grossi (Utrecht) addressing the relationship between collective and individual obligations in a framework based on dynamic deontic logic. After the afternoon coffee-break, Franco Raimondi (King's College London) presented a paper on the specification and verification of the correct functioning behaviour of multi-agent systems and the epistemic properties of a system by means of model checking via ordered binary decision diagrams. The presentation by Olga Pacheco (Minho) on issues of delegation in a role-based organisation concluded the day.

The second day began with an invited talk by Mark Brown (Syracuse) on the technical and philosophical issues that arise in deontic logic from attempts to model obligations, contracts, and negotiation. This was followed by Wojciech Jamroga (Twente) reporting on a combination of ATL and standard deontic logic, thereby allowing the representation of abilities as well as obligations. The third presentation of the day was given by Robert Demolombe (Toulouse) who spoke about work on the situation calculus and dependence logic to model changes in obligations as a consequence of actions. He was followed by Lennart Åqvist (Uppsala) who presented soundness and completeness results for particular classes of combinations of tense and deontic modalities. The last talk of the day, by Andrei Kouznetsov (Kemerovo), concerned the formalisation of weak deontic logics on quasi-matrix semantics.

The first session on the 28th of July started with Antonino Rotolo (Bologna) reporting on a deontic defeasible logic for modelling the concepts of intention and obligation. This was followed by a presentation by Jorg Hansens (Leipzig) who reported on metalogical results on logics for conflicting imperatives represented by dyadic deontic operators. After lunch, Jan Broersen (Utrecht) presented a dyadic deontic logic with which to model deadlines for a group of agents and its translation into standard branching time temporal logic. After lunch Peter Vranas (Iowa) introduced a new approach to imperative logic. Adam Wyner (King's College London) gave the last talk before the afternoon coffee-break; his talk was on the use of deontic logic for the representation of obligations in the context of contract formation, execution, and monitoring. The last talk of the workshop, by Lou Goble (Willamette), was on soundness, completeness, and decidability of logics for deontic dilemmas. A lively discussion between a panel and the audience on the trends and open problems on deontic logic in computer science preceded the social dinner. Marek Sergot (Imperial College) was prompted to share, in a post-dinner speech, his recollections of the highlights of each DEON workshop since the first meeting more than 10 years ago. Not all that was reported seemed to be entirely academic in nature.

The proceedings of the workshop were published by Springer as LNCS volume 3065 and distributed at the workshop. Authors of a number of selected of papers will be subsequently invited to submit articles to a special issue of the Journal of Applied Logic, published by Elsevier.

Alessio Lomuscio (University College London), and Donald Nute (Georgia) chaired a team of 23 reviewers from 11 countries. The workshop was locally organised by José Carmo (Madeira), Eduardo Fermé (Madeira), and Filipe Santos (ISCTE Portugal).

DEON04 received financial support from FCT — Fundação para a Ciência e Tecnologia, FLAD — Fundação Luso-Americana para o Desenvolvimento, CITMA — Centro de Ciência e Tecnologia da Madeira, the Journal of Applied Logic, as well as AgentLink III.

Symposium on Adaptive Agents and Multi-Agent Systems

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

The Symposium on Adaptive Agents and Multi-Agent Systems took place in April 2004 in Leeds, UK. Now in its fourth edition, the Symposium is well-established event with a sizeable group of regulars, but having also many new participants.

The Symposium brings together experts in Machine Learning and Agent-Based Systems, as well as the combination of both. This year, it featured 18 papers by international researchers from as close as York, UK, and as far as Kyoto, Japan. The research areas covered were just as diverse — from common (and uncommon) reinforcement learning, through learning of reputation, to evolution of cellular automata — as expected from such broad and multi-disciplinary research community.

This year's keynote address, thanks to sponsorship by AgentLink III, was given by Jürgen Schmidhuber on the subject of "Gödel Machines and Other Wonders of the New, Rigorous, Universal AI". As the title itself suggests, it was a highly inspiring and stimulating talk.

As usual for the Symposium, much time was reserved for the participants to chat and discuss their work over lunch and coffee. Of course, there was also the (by now obligatory) Symposium dinner with good food (the best of English cuisine, also known as Indian), drinks, and entertaining conversations.

The Symposium proceedings are available from AISB (www.aisb.org.uk). For those who are patient enough to wait, a Springer LNAI volume will be published later this year featuring the best papers of the last two Symposia.

At the concluding discussion, the delegates agreed that this event should continue in the future. One major change is that next year it will be split off from AISB (under whose umbrella the first four Symposia took place). The next Symposium, to take place in April 2006, in Paris, will be a stand-alone event. For up-to-date information, look out for the upcoming call for papers and of course the Symposia webpage at http://www.aamas.net.

European Workshop on Multi-Agent Systems

Mark d'Inverno
University of Westminster
United Kingdom
M.dInverno@westminster.ac.uk

Carles Sierra
IIIA-CSIC
Spain
sierra@iiia.csic.es

Franco Zambonelli
Università di Modena e Reggio Emilia
Italy
franco.zambonelli@unimo.it

The 1st European Workshop on Multi-Agent Systems (EUMAS 03) was held at St. Catherine's College, Oxford University, over two days in December 2003. This workshop was organised in direct response to the increasing number of European researchers and projects in the multi-agent systems field and to the ever-expanding AgentLink initiative, the EU-funded network of excellence in the area.

It was clear to us and others in the field that there was a need for a regular opportunity for the European multi-agent system community to get together. Not only did we want such a regular forum in order to exchange research results and ideas, strengthen research collaborations, and so on in a structured workshop, but also the chance to get together over more informal settings in the workshop bar to catch up on what we all had been up to in the last year — and of course often the best ideas come at the bar! By keeping costs as low as possible, we attracted many PhD students, and one of the most significant benefits of the workshop was that it gave PhD students direct access to some of the leading agent researchers whose papers were shaping their own PhD work.

The location of St. Catherine's College, Oxford University (where one of the chairs had enjoyed life as an undergraduate for several years) was chosen because it had excellent facilities for running a small workshop. In addition, it had reasonable travel connections to the rest of the continent, and had also been used very successfully on previous occasions to run the smaller UK Multi-Agent Systems Workshop (UKMAS 2000 and UKMAS 2001).

This workshop was organised in direct response to the increasing number of European researchers and projects in the multi-agent systems field and to the ever-expanding AgentLink initiative, the EU-funded network of excellence in the area.

It was very important to us, as the general and programme chairs, that the workshop was an *informal* forum and therefore a workshop in the true sense of the word rather than as just another mini-conference. The problem with some such events is that the publication, which did not necessarily have the widest uptake, becomes more important than the event itself. We decided therefore not to provide a formal publication and in this respect we made it clear, in our submission guidelines, that we welcomed both recently published as much as previously unpublished material. Of course not having a formal publication was a concern for us in terms of the number of people wishing to attend; we were only too aware of the demands on many academics to keep publishing, and moreover that travel money was often only ever available with the promise of a resulting publication. We need not have worried too much as the call for contributions attracted more than 120 submissions, a large percentage of which contained original research. From these submissions the organisers selected around 60 papers for presentation and the workshop eventually attracted around 130 delegates.

The first EUMAS was a very intensive workshop, yet friendly and informal, and the programme included two invited speakers. The first was Wiebe van der Hoek from the University of Liverpool; the talk was on *Reasoning about Knowledge, Rationality and Action*, which covered epistemic logics and the design of knowledge-based communication protocols. What was excellent about this presentation was that it showed how it is possible to give a talk about formal structures and logic in an accessible and entertaining way. The second was by Michael Luck who gave an engaging account of the road-

map for multi-agent systems research that he had recently co-authored as part of his role as AgentLink II coordinator. This is available online at http://www.agentlink.org, and is called Agent Technology: Towards Next Generation Computing. In addition, there were two panel sessions, the first entitled Future Scenarios and Potential Dangers of Agent-Based Systems and the second called Challenges in Agent-based Modelling. The paper presentations were organised into 8 sessions that, because of the large number of acceptances, ran in parallel.

Thanks were largely due to Michael Luck and Steve Willmott, on whose experience we drew on an almost daily basis, and to Serena Raffin for making the whole thing run so smoothly. As a result of the success of the event, and in recognition of the high-quality research that was presented, the workshop chairs have been asked to select what they considered to be some of the best papers from this event for a special issue of the Applied Artificial Intelligence Journal, which will be appearing next year. Most importantly though, and on behalf of the EUMAS organising committee for 2004, we would like to invite you to attend the Second European Workshop on Multi-agent Systems, which is to be held later this year in Barcelona. Please visit www.eumas.org for more details.

Symposium on Logic-Based Agent Verification

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

The Symposium on Logic-Based Agent Verification was the second event within the "Logic and Agents" activity in CoLogNET, which is also partially supported by AgentLink. Further details can be found at:
http://www.csc.liv.ac.uk/~michael/symposium04.CALL/

Background

The main aim of this activity, on "Logic and Agents", is to bring together expertise from the Computational Logic and Agent-BasedSystems communities, via interaction between the CoLogNET and AgentLink European coordination actions, respectively. Formally, this activity is one of the main areas within CoLogNET and was previously a special interest group within AgentLink.

The activity is funded by both actions, showing their commitment to this important area. This formalised activity aims to allow expertise/ideas to flow between the areas, thus providing greater chance of productive technology transfer. Therefore, in the future, this activity will contribute to the uptake of logic-based agent technology and logical methods for agent-based systems. This, in turn, may impact on forthcoming standards, for example where semantics/verification/compliance is an issue.

CoLogNET

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

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

Among its many activities, CoLogNET supports the "Logic and Agents" activity, associating with AgentLink. This activity is coordinated by Michael Fisher.

Symposium

The work of this "Logic and Agents" activity is centred around two symposia, one on "logic-based agent implementation", the other on "logic-based agent verification", both of which provide very strong links between the actions. The first of these was held on 3rd February 2003 in Barcelona; the second was held on 9th July 2004 in Liverpool.

The aim of this second symposium, on "logic-based agent verification", was to consolidate expertise, stimulate research/collaboration in each area, showcase successful applications, and enhance the uptake of logic-based agent technology.

This symposium was co-located with a large event within the agent community: the Sixth European Agent Systems Summer School (EASSS – http://www.agentlink.org/happenings/easss/2004). The original aim was that each symposium should be co-organised by a leading EU researcher in the topic area. We are pleased that such a researcher agreed to be involved in this way, and we thank Mike Wooldridge (Univ. Liverpool, UK), who co-organised the symposium on logic-based agent verification.

We would like to acknowledge administrative support from the AgentLink local organisers. Funding for this event was provided by both CoLogNET and AgentLink, thus we were able to support a number of speakers, allowing them to participate in the event. We were also able to attract a high-profile non-EU keynote speaker from Australia.

Summary

The presentations not only gave a broad view of the state of the art in logic-based agent verification throughout both communities, but also allowed dissemination of key breakthroughs in this area. We believe that this event will enhance collaboration between groups, projects, academia/industry, etc.

Finally, not only did this symposium allow participants to showcase successful applications concerning the verification of agents using logic-based approaches, but it also provided an opportunity to identify important future research issues and key requirements for wider use of these techniques.

A special issue of
the Journal of Applied Logic
on the topic of
Logic-Based Agent Verification
is open for submissions
until **15th October 2004**

For details, see:
http://www.csc.liv.ac.uk/~michael/LBAV04

Developing Intelligent Agent Systems

A Review by
Terry R. Payne
University of Southampton
United Kingdom
trp@ecs.soton.ac.uk

Lin Padgham
Michael Winikoff

ISBN 0-470-86120-7

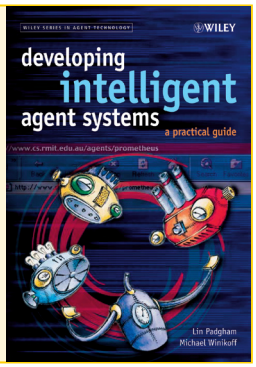

After at least a decade of Agent-Based Research, agents are no longer simply abstract artefacts or theoretical notions shared between AI practitioners, but rather have emerged as well understood, deployed entities appearing within a plethora of Multi-Agent Systems (within industrial as well as academic contexts). Since the early 1990's, many facets of agent research have received attention: from communication metaphors, agent communities (and then corresponding architectures), negotiation, and reasoning mechanisms, to agent architectures and agent-oriented programming languages. Research into the latter has resulted in the emergence of a variety of platforms and tools that facilitate the creation of agent systems. This includes the specification of individual plans, architectural designs (such as BDI), desired emergent behaviour, and the resultant roles that different classes of agents may assume. However, for there to be a growth in the successful transition of such technology from academia to industry, it is critical that well understood design methodologies emerge for agent-based systems, and that such methodologies are accessible to students and IT professionals in a clear and concise way.

The Prometheus methodology, described in detail in Padgham and Winkoff's book "Developing Intelligent Agent Systems", is one of the most mature methodologies currently proposed. Developed since the mid 90's, it consists of a detailed process for specifying, designing, implementing and testing/debugging agent-oriented software systems. Using both graphical and structural elements, a top-down design framework is provided, by initially considering high-level concepts, such as goals, plans (that achieve the goals), the necessary percepts and actions, and different scenarios in which the goals can be achieved. Additional artefacts can be constructed (through successive iterations and refinements of the design) which result in a transition through three design phases: *the System Specification, Architectural Design and Detailed Design* phase.

- The **System Specification Phase** focuses on identifying the goals and basic functionalities of the system, along with the inputs (percepts) and outputs (actions).
- The **Architectural Design Phase** uses the outputs from the previous phase to determine which agent types the system will contain and how they will interact.
- The **Detailed Design Phase** looks at the internals of each agent and how it will accomplish its tasks within the overall system.

This book, aimed primarily at industrial software developers and undergraduate students, takes a pragmatic, and easy to follow journey through this design process, grounded by the development of a case study (an "Electronic BookStore", listed in detail in Appendix 1). Whilst some notions of agency are introduced within the first two chapters, the book avoids dwelling on the different facets of the Agent research field, and rather focuses directly on the elements necessary for designing an agent system. Equally, the book remains essentially agnostic with regards to an implementation framework, describing a design mechanism that could be applied to various platforms.

The book is divided into five main parts, and includes appendices that provide additional detail on some of the notation used, as well as a full listing of the descriptor forms used in the book. The first part (consisting of chapters 1 to 3) introduces the reader to the concept of *agents*, briefly justifying the use of agents for different classes of problems, and then expands on agent-oriented design methodologies, by introducing the different concepts necessary for building agent systems. Prometheus and its principal design stages are then introduced, and contrasted to other agent-oriented methodologies.

The second part (chapter 4) commences the design process with a high-level specification of *Goals, Functionalities, Scenarios, and Interfaces* that constitute the System Specification phase. Each of the system-level artefacts are presented, and the constituent fields within their descriptor forms are described. Approaches for refining the design are then described, followed by heuristics for checking completeness and consistency of the design.

The third part (chapters 5 to 7) focuses on the Architectural Design Phase, consisting primarily of the decision on the types of agents that would be used in the system, understanding the necessary interactions between agents, and then developing the overall system structure. The fourth part describes the Detailed Design Phase, where the agents themselves are decomposed into specifications of their capabilities and processes. The final part (chapter 10) preserves the pragmatism of the whole book, by illustrating how the resulting design could be implemented on an existing agent platform; the one used here is JACK (developed by the company Agent Oriented Software), though, as the authors point out, there are several other agent platforms that could also be used. Although little code is actually given (it would be unreasonable to include code that actually implemented the resulting capabilities within the book), this chapter conveniently summarises the significant artefacts generated by the design, and could easily be used as a guide for mapping these artefacts to other agent platforms.

Throughout the book, emphasis is placed on the designer fully specifying each of the artefacts and their descriptor forms at each part of the design process, with checks for completeness and consistency. Useful tips and "rules-of-thumb" are presented to simplify the process, and to identify possible pitfalls or problems. Examples are used in each of the phases to illustrate the design process, in addition to the bookstore case study which is developed throughout the book. Further, the Prometheus design tool can be obtained from http://www.cs.rmit.edu.au/agents/Prometheus.

Despite the fact that a detailed bibliography is provided, "*Developing Intelligent Agent Systems*" is somewhat light on the theory behind agent-oriented programming and system design, and avoids details regarding architectural components of large-scale agent systems (such as issues to do with discovery, mobility, etc.). Likewise, little is said about issues pertaining to heterogeneous, or open multi-agent systems. Although heterogeneity is of interest to many researchers, most agent deployments are designed to achieve specific goals and are typically implemented using a single agent development platform. Thus, the limitations of this book are actually its strengths, as IT developers can pragmatically focus on rapidly designing and deploying a working multi-agent system.

To conclude, the authors have done an excellent job in describing the different stages involved in designing agent artefacts with Prometheus, and clearly demonstrate how these can be used to design, analyse, and refine multi-agent systems by means of a detailed running case study. This book is a valuable contribution, not only in terms of teaching agent-oriented software engineering and considering design decisions associated with developing a multi-agent system, but in the pragmatic sense of "how are agents actually built". For Agent-based research to have an impact on IT and software development, agent principles and their use should be clearly accessible to IT professionals charged with the task of implementing agent systems. This book goes a long way towards achieving this goal.

Cybele: An Agent Infrastructure for Modelling, Simulation, and Decision Support

Vikram Manikonda, Goutam Satapathy and Renato Levy
Intelligent Automation Incorporated
USA
{vikram, goutam, rlevy}@i-a-i.com

Over the last decade, there has been an increasing number of commercial and defence programs, in the United States and abroad, that have identified the need for agent technology in applications involving modelling, simulation, and decision support. Several initial prototypes have been developed and are currently being transitioned into products. Application areas include air and ground transportation networks, military logistics, networking, and robotics [1,2,3]. To meet the needs of this growing market, over the last several years Intelligent Automation Inc. (IAI) has developed Cybele®, a robust and efficient agent infrastructure for the development and deployment of large scale distributed agent applications.

Architecture

Cybele© is a runtime environment, built on the top of the Java 2 platform, for the management and execution of agents. The architecture consists of a kernel and several service implementations. The architecture kernel provides certain application interface methods for agent programmers to write classes representing activities using the Activity Centric Programming (ACP) paradigm [4]. These methods are called Activity Oriented Programming Interface (AOPI) methods. The AOPIs in turn use the published interfaces to different agent services. While the execution environment supporting ACP is a layer above a Java Virtual Machine, the activity classes also use java packages in addition to AOPI (see Figure 1).

The agent infrastructure adopts a service-layered architecture promoting plug-and-play capability of agent services (see Figure 2). The services and their interfaces are defined in such a way that performance can be fine-tuned by loading different service implementations as appropriate to the operating system/platform/network and/or the agent application domain, without having to rewrite the agent code. Agent services are categorised into three types of layers — basic, fundamental, and supplemental, based on their relevance to a typical agent runtime environment. We categorise basic services as essential services required for creating agents, activities, agent operation, and inter-activity communication. They include error handling, concurrency management, event handling, thread management, and internal event services. The fundamental services are primarily event generation services that are essential for autonomous and communicative multi-agent system. Services included in the fundamental layer are communication, timer, data sharing, and GUI services. The supplemental services are services that enhance performance; these include migration and load balancing services.

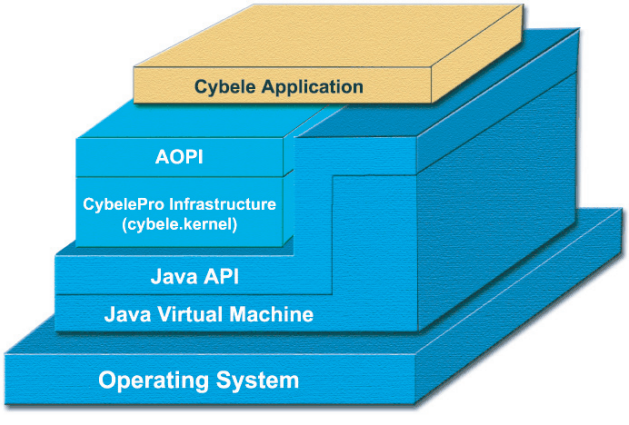

Figure 1: Cybele runtime environment

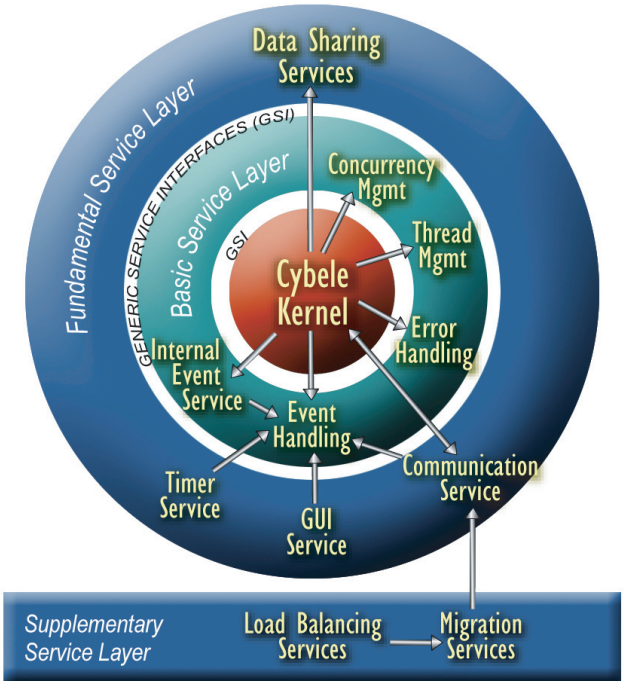

Figure 2: Cybele Service Layered Architecture

Features

In Cybele, agents are autonomous and event driven, and do not share data. Events currently supported include message, timer, and agent internal events. Cybele encourages an ACP approach and provides support for different concurrency models between activities of an agent as well as data sharing among activities for high efficiency. Location-independent communication between agents is supported via publish-subscribe based messaging, with support for synchronous, asynchronous, broadcast, multicast, and point-to-point messaging. Both continuous and discrete clock capabilities (for event driven large-scale distributed system) are supported by Cybele. The infrastructure also provides users with a remote internet browser-based agent management tool for location independent launching of Cybele agent containers (i.e., the agent runtime environment), and a directory service (yellow-page and white-page) that is useful for developing open systems. FIPA compliance is also supported via Cybele-FIPA, a FIPA compliant agent platform that has been implemented as a layer on top of Cybele. Cybele-FIPA offers Agent Management System and Directory Facilitator FIPA services and a HTTP-XML gateway based interoperability solution.

Cybele has been customised to meet specific customer and user community requirements, and currently three implementations of Cybele exist. These include OpenCybele, an open source version, CybelePro, a commercial version, and CybeleHLA, a version of Cybele that is compliant with the US Department of Defense's High Level Architecture (HLA). CybeleHLA was developed for the US Department of Defense and NASA. We refer the reader to www.cybelepro.com for more information on these infrastructures.

Applications

Over the last few years, Cybele has been used in several applications such as modelling and simulation of the US National Airspace (NAS), military logistics and battle management, agent-based scheduling systems, collaborative decision support systems, and ad-hoc network simulations. Two of these examples are discussed next.

Figure 3 shows the Airspace Concept Evaluation System (ACES) [5,6] that is being developed using Cybele. ACES is a fast-time agent-based system which provides researchers with the capability to model, simulate, and assess the per-

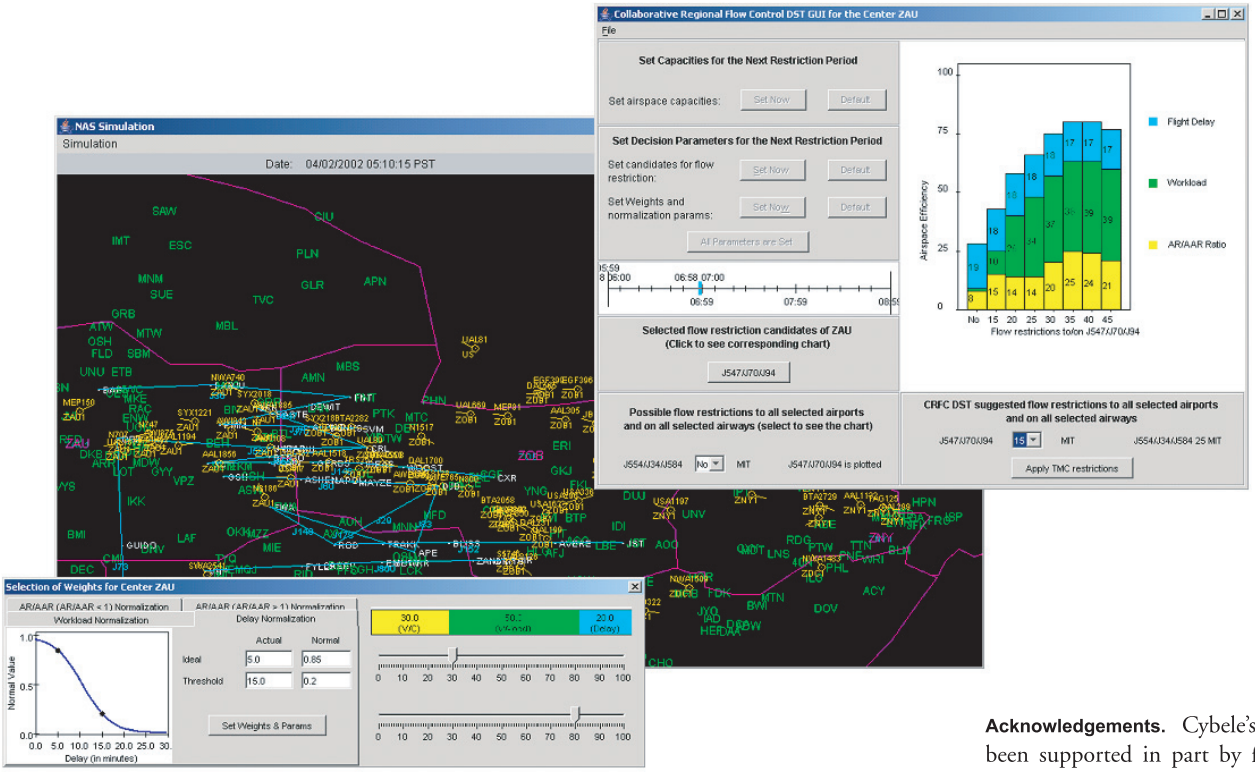

Figure 4: Decision Support Tool for Collaborative Regional Flow Control

formance of the NAS under current and future air traffic management concepts[1] . At the core of the system is the Cybele™ agent infrastructure integrated with the DOD High Level Architecture (HLA). Agents in the simulation represent NAS elements such as flights, airports, terminal areas, controllers, and airline operation centres with varying levels of fidelity. An individual simulation is made up of a collection of tens of thousands dynamically interacting NAS agents configured to represents a concept/scenario. The Cybele-HLA architecture supports repeatable, distributed, discrete event-based simulation capability, run-time centralised simulation monitoring, centralised logging capability for debugging, decentralised local data collection and centralised post-processing, centralised simulation/federation management with multiple run capability, and centralised random number generation capability. A typical ACES simulation involves approximately 38,000 agents, 150,000 activities, and approximately 3 million messages exchanged across 300,000 communication channels. One such simulations represents a day of NAS activities (28 hours of simulated time) and runs in about 12 hours clock time on 10 high-end PC's (2.4 GHz, 2GB memory machines).

Figure 4 shows a Distributed Decision Support Tool (DST) developed for Collaborative Regional Flow Control (CRFC) using Cybele. CRFC is a new air-traffic management concept for optimising airspace efficiency, in which the traffic management coordinators (TMC) collaboratively determine optimal flow restrictions across regional boundaries [7]. The concept is implemented as a distributed agent-based collaborative DST that evaluates projected airspace efficiency measures as a trade-off between TMC set decision factors. DST uses a Decision Support System Infrastructure (DSSI) and its game-theoretic libraries to perform an iterative negotiated search and optimisation to arrive at collectively agreeable flow restrictions (local equilibrium). The DSSI is built on top of the Cybele™ and includes agent decision-making models such as state machines, Markov Decision Process (MDP), Partially Observable MDP, and interaction and negotiation models for communicating decision-making agents.

Figure 3: ACES Simulation Testbed

1. ACES is being developed by a Raytheon Led team for NASA Ames under the NASA's Air Traffic Management-System Development and Integration effort.

Acknowledgements. Cybele's development has been supported in part by funds from the US Department of Defense, Department of Commerce, NASA, and IR&D. The authors would like to thank these agencies for their support; we also thank the IAI Cybele development team for their technical contributions.

References

1. Defense Advanced Research Program Agency. See www.darpa.mil.
2. NASA Virtual Airspace Modeling and Simulation Concept (VAMS). See www.vams.arc.nasa.gov.
3. NASA Human and Robotic Technology Initiative. See www.exploration.nasa.gov .
4. K. Erol, J. Lang and R. Levy. Designing Agents from Reusable Components. In proceedings of the Fourth International Conference on Autonomous Agents, Barcelona, Spain 2000.
5. D. Sweet, V. Manikonda, J. Aronson, and Karlin Roth. Fast-Time NAS Simulation for Analysis of Advanced ATM concepts. In AIAA Modeling and Simulation Conference and Exhibit, 5 – 8 Aug 2002 Monterey, California.
6. J. Aronson, V. Manikonda, W. Peng, R. Levy and K. Roth. An HLA Compliant Agent-based Fast-Time Simulation Architecture for Analysis of Civil Aviation Concepts. Spring SISO Simulation Interoperability Workshop, April 2003, Orlando, Florida.
7. G. Satapathy, V. Manikonda, J. Robinson, and T. Farley. En-Route Sector Metering using a Game-theoretic Approach. In workshop on Game Theoretic and Decision Theoretic Agents, Eighteenth National Conference on Artificial Intelligence, Edmonton, Alberta, Canada, July 28 – Aug 1, 2002.

AGENT EVENTS CALENDAR

For further details visit:

www.agentlink.org/happenings/other-events.php

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

2004			
WITSE 2004	Sep 20–25	2nd Workshop on Intelligent Technologies for Software Engineering	Linz, Austria
ACSys 2004	Sep 27	Workshop on Agents for Complex Systems	Timisoara, Romania
CIA 2004	Sep 27–29	Cooperative Information Agents	Erfurt, Germany
JELIA 2004	Sep 27–30	9th European Conference on Logics in Artificial Intelligence	Lisbon, Portugal
GSEM 2004	Sep 27–30	The International Conference on Grid Services Engineering and Management	Erfurt, Germany
COMPSAC 2004	Sep 28–30	28th IEEE Annual International Computer Software and Applications Conference	Hong Kong, China
SBIA 2004	Sep 28–Oct 1	XVII Brazilian Symposium on Artificial Intelligence	São Luis, Maranhão, Brazil
MATES 2004	Sep 29–30	2nd International German Conference on Multi-agent System Technologies	Erfurt, Germany
CLIMA V	Sep 29–30	5th International Workshop on Computational Logic in Multi-Agent Systems	Lisbon, Portugal
ABA 2004	Oct 10–13	The 2004 International Conference-Crossing Workshop on Autonomous Computing and Agents for Business Automation	The Hague, The Netherlands
ESAW 2004	Oct 20–22	5th International Workshop Engineering Societies in the Agents World	Toulouse, France
AAC-GEVO 2004	Oct 21–24	International Workshop on Agents and Autonomic Computing, and Grid Enabled Virtual Organizations	Wuhan, China
KSCO 2004	Oct 26–28	3rd International Conference on Knowledge Systems for Coalition Operations	Pensacola, FL, USA
ICDM 2004	Nov 1–4	4th IEEE International Conference on Data Mining	Brighton, UK
FOIS 2004	Nov 4–6	Formal Ontology in Information Systems	Torino, Italy
ISWC 2004	Nov 7–11	3rd International Semantic Web Conference	Hiroshima, Japan
EUSAI 2004	Nov 8-10	2nd European Symposium on Ambient Intelligence	Eindhoven, The Netherlands
CoMoA 2004	Nov 8–12	Workshop on COnceptual MOdelling for Agents	Shanghai, China
ICTAI 2004	Nov 15–17	16th IEEE International Conference on Tools with Artificial Intelligence	Boca Raton, USA
AISTA 2004	Nov 15–18	International Conference on Advances in Intelligent Systems - Theory and Appliciations	Kirchberg, Luxembourg
IBERAMIA 2004	Nov 22–26	9th Ibero-American Conference on AI	Puebla, Mexico
EU-LAT	Nov 30–Dec 3	Workshop in e-Environment	San Jose', Costa Rica
HIS 2004	Dec 5–8	4th International Conference on Hybrid Intelligent Systems	Kitakyushu, Japan
IWMMAS 2004	Dec 10–11	International Workshop on Massively Multi-Agent Systems	Kyoto, Japan
EUMAS 2004	Dec 16–17	2nd European Workshop on Multi-Agent Systems	Barcelona, Spain
ICMLA 2004	Dec 16–18	The 2004 International Conference on Machine Learning and Applications	Louisville, USA
2005			
SAINT 2005	Jan 31–Feb 4	The 2005 International Symposium on Applications and the Internet	Trento, Italy
WM 2005	Apr 11–13	Workshop on Peer-to-Peer and Agent Infrastructures for Knowledge Management	Kaiserslautern, Germany
KIMAS 2005	Apr 18–21	International Conference on Integration of Knowledge Intensive Multi-Agent Systems	Waltham, USA
FLAIRS 2005	May 16–18	18th International FLAIRS Conference	Clearwater Beach, USA
AAMAS 2005	July 25–29	4th International Joint Conference on Autonomous Agents and Multi-agent Systems	Utrecht, The Netherlands
IJCAI 2005	July 30– Aug 5	19th International Joint Conference on Artificial Intelligence	Edinburgh, UK
CEEMAS 2005	Sep 12–14	4th International Central and Eastern European Conference on Multi-Agent Systems	Budapest, Hungary
BV&AI 2005	Oct 19–21	1st International Symposium on Brain, Vision and Artificial intelligence	Naples, Italy

About AgentLink...

AgentLink is the European Coordination Action for Agent-based Computing, a network of researchers and developers with a common interest in agent technology. It is funded by the European Commission, so membership is free, but you can only take advantage of AgentLink activities if you are a member. If your organisation is engaged in agent-related activities, you should join. Contact the coordinator or administrator for details, or visit www.agentlink.org.

The aim of AgentLink News is to provide an informal way of communicating both what is happening in AgentLink and the agent world generally.

AgentLink News offers a range of articles including features, reports on conferences and workshops, informal description of research results and new software, book reviews, and website developments.

In order to provide the most relevant news, we invite you to contribute. If you are interested in writing an article for the newsletter, please contact our editor, Rafael H. Bordini, at the address below. The next closing date for the submission of articles is 31st of October, 2004.

AGENTLINK NEWS EDITOR
Rafael H. Bordini
Department of Computer Science
University of Durham
Durham DH1 3LE
United Kingdom
[e] editor@agentlink.org

AGENTLINK PUBLICATIONS AND WEB COORDINATOR
Serena Raffin
School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ
United Kingdom
[e] web@agentlink.org

AGENTLINK ADMINISTRATOR
Adele Maggs
Department of Computer Science
University of Liverpool
Liverpool L69 3BX
United Kingdom
[e] admin@agentlink.org

Rebecca Earl
School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ
United Kingdom
[e] publications@agentlink.org

AGENTLINK EVENT COORDINATOR
Catherine Atherton
Department of Computer Science
University of Liverpool
Liverpool L69 3BX
United Kingdom
[e] events@agentlink.org

AGENTLINK COORDINATORS
Peter McBurney
Department of Computer Science
University of Liverpool
Liverpool L69 3BX
United Kingdom
[e] Peter.McBurney@agentlink.org

Terry Payne
School of Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
United Kingdom
[e] Terry.Payne@agentlink.org

AGENTLINK EXECUTIVE COORDINATORS
Michael Luck, University of Southampton, UK
Michael Wooldridge, University of Liverpool, UK

INDUSTRY ACTION (Workpackage 1) COORDINATOR
Michal Pechoucek
Czech Technical University
Czech Republic
[e] Michal.Pechoucek@agentlink.org

STANDARDISATION ACTIVITY (Workpackage 2) COORDINATOR
Monique Calisti
Whitestein Technologies AG
Switzerland
[e] Monique.Calisti@agentlink.org

RESEARCH ACTION (Workpackage 3) COORDINATOR
Steven Willmott
LSI,Universitat Politècnica de Catalunya
Spain
[e] Steven.Willmott@agentlink.org

STUDENT INTEGRATION PROGRAMME (Workpackage 4) COORDINATOR
Wiebe van der Hoek
University of Liverpool
United Kingdom
[e] Wiebe.van.der.Hoek@agentlink.org

TECHNICAL FORUM (Workpackage 5) COORDINATOR
Andrea Omicini
DEIS, Università di Bologna
Italy
[e] Andrea.Omicini@agentlink.org

TECHNOLOGICAL ROADMAP (Workpackage 6) COORDINATOR
Michael Luck
University of Southampton
United Kingdom
[e] Michael.Luck@agentlink.org

ADDITIONAL COMMITTEE MEMBERS
Onn Shehory
IBM Research
Israel
[e] Onn.Shehory@agentlink.org

Simon Thompson
BT Exact
United Kingdom
[e] Simon.Thompson@agentlink.org

Jörg Müller
Siemens AG
Munich
[e] Joerg.Mueller@agentlink.org

ISSN 1465-3842