

Agents on the web

ANYONE who's ever arrived in an unfamiliar city and wished for a personal guide to help them make the most of their visit will appreciate the vision behind the Agentcities project.

Funded by the EU 5th Framework, Agentcities is aiming to deploy intelligent network services capable of advising mobile users about an environment in which they find themselves.

One Agentcities scenario envisages planning an evening's entertainment for a traveller, selecting dining, entertainment and accommodation according to their preferences.

Interface

Agentcities is funding small-scale projects to deploy existing services on the emerging Agentcities network. In the second half of 2002, the Aberdeen AKT team has been working to interface selected AKTbus components to the Agentcities infrastructure.

This work will allow AKT knowledge services to operate as software agents and serve Agentcities users. In conjunction with the Aberdeen software agents group, a demonstrator is being developed that offers entertainment options to people visiting Scotland's Granite City.

Next-generation

THE World Wide Web is like a large library. But searching through it is unintelligent and often painful. As it gets larger, the dangers of infomog will increase.

The Next-generation Semantic Web is intended to get around these problems by developing ways of sifting intelligently through information. By providing different layers of representation for content, the Semantic Web allows increasingly sophisticated and more powerful inferences to be made about the content, as the details are abstracted from the raw material.

For instance, in our figure (opposite), we see an actual news story as it appears in the Open University's KMi Planet e-journal, with a headline 'Outstanding contribution award for Daisy Mwanza'. Here, the content has been 'rendered'

– generated on the fly – for display on the web. The layer above shows the HTML encoding, which gives the 'look' of the page behind the scenes. For instance, we can see that what makes the headline the headline is the pair of tags, <H1> and </H1>, which show where it begins and ends.

Abstract


The next layer above is the extended mark-up language XML, which allows a more abstract – and hence more powerful – representation. Now the headline is delimited by the tags <Headline> and </Headline>. These tags are arbitrary and their only meaning is by mutual agreement by specific communities of users. So the people behind KMi Planet have decided that <Headline> should be interpreted as <H1>. But mobile phone users, for example, may

interpret the <Headline> tag differently to allow it to appear in a very small font. Hence XML allows what is called repurposing of content.

So far so good. But we want intelligent look-up. If we move up another layer, we can now see some inferential capability, via the **Resource Description Framework (RDF)**. RDF represents relations as triples, two objects and a relation between them. Both objects and the relation are actually locations on the web.

This is a simple but amazingly powerful idea, allowing much more complex searches than are possible on the World Wide Web, enabling principled, large scale indexing and retrieval by machines.

But it is not yet intelligent. The top layer of representation shows a language called **OWL (Web Ontology Language)**. This, and languages like it, provide the capability to reason about the relations stored in RDF triples, thus going



Better by design

AKT is strongly concerned with industrial relevance. As one example of this, we are working closely with Rolls-Royce to provide support for its aero engine designers. In collaboration with the Transmissions and Structures team at Derby, work is under way on developing a Designers' Workbench software system.

The current prototype allows designers to interact with a familiar engineering drawing interface. As their design progresses, the Workbench checks that the new parameters are in accordance with a variety of pre-existing constraints, using constraint-satisfaction technology developed by the Aberdeen AKT group.

Workbench users can also access the rationales for these constraints and should they wish to override them, they can provide their own rationale for doing so. In this way, the Workbench builds up a repository of design knowledge and experience.

The prototype Designers' Workbench will be trialled by Rolls-Royce early in 2003.

Going to town on technology

ALTHOUGH AKT has a remit from the EPSRC to think as broadly and as innovatively as possible, this does not mean we spend all our time in the 'blue sky'. Relevance to the real, information-related problems in industry and within organisations that follow from changes in technology is a priority.

As part of our commitment to industrial relevance, AKT has instituted a series of 'town meetings', in which our technologies can be seen in action, via presentations, demonstrations and informal chats – those 'coffee break' moments where the serious questions can be asked.

The first of these meetings was held on September 18, 2001, at the Institute of Materials in London. The AKT Manifesto was launched (<http://www.aktors.org/publications/Manifesto.doc>) and technology was demonstrated, including AKT-0 (an early integration of AKT technologies), tools and

methods for enriching web-based knowledge content and language-processing technology.

An exciting day was enjoyed by all. The audience included representatives of a number of enterprises ranging from massive bluechips to the nimblest technology startups.

Opportunity

A second meeting is in order, to demonstrate the continuing success of AKT and to showcase the achievements of the last 12 months. It will provide an opportunity for individual organisations to discuss specific knowledge management challenges and how these relate to AKT technologies.

Places are limited for this one-day event, and are by invitation only. If you have an interest in attending, please email Susan Davies (AKT administrative co-ordinator) at sdd@ecs.soton.ac.uk, or telephone 023 8059 3523.

<http://www.aktors.org>: it's no ordinary website

A WEBSITE is essential nowadays for scientific, technological or business communication and, in common with most contemporary collaborative projects, AKT has a website which can be found at the address above.

But [aktors.org](http://www.aktors.org) is no ordinary website – we have applied our own principles and technologies for the creation of the site.

The site contains information about AKT and its aims, associated events and the AKTors Club. There are links to the AKT Manifesto and to the individual websites of the consortium partners.

But site editor Hugh Glaser is most excited about the use of our own technology in the site. Much of our internal knowledge management is done via the project Intranet, but even on the

public site we can glimpse the future – a website incorporating ontologies, authoring tools and instantly generated discussion environments.

The welcome page includes the latest story from *AKT Planet*, an intelligent online news server and our ontology-mediated e-print archive gives immediate access to the technical papers produced within the project.

There is also access to public discussion spaces of the papers, using the Digital Document Discourse Environment (D3E).

We are committed to providing the best technology for information dissemination. The beauty of it is that we can use it to tell the world about ourselves!

Semantic Web

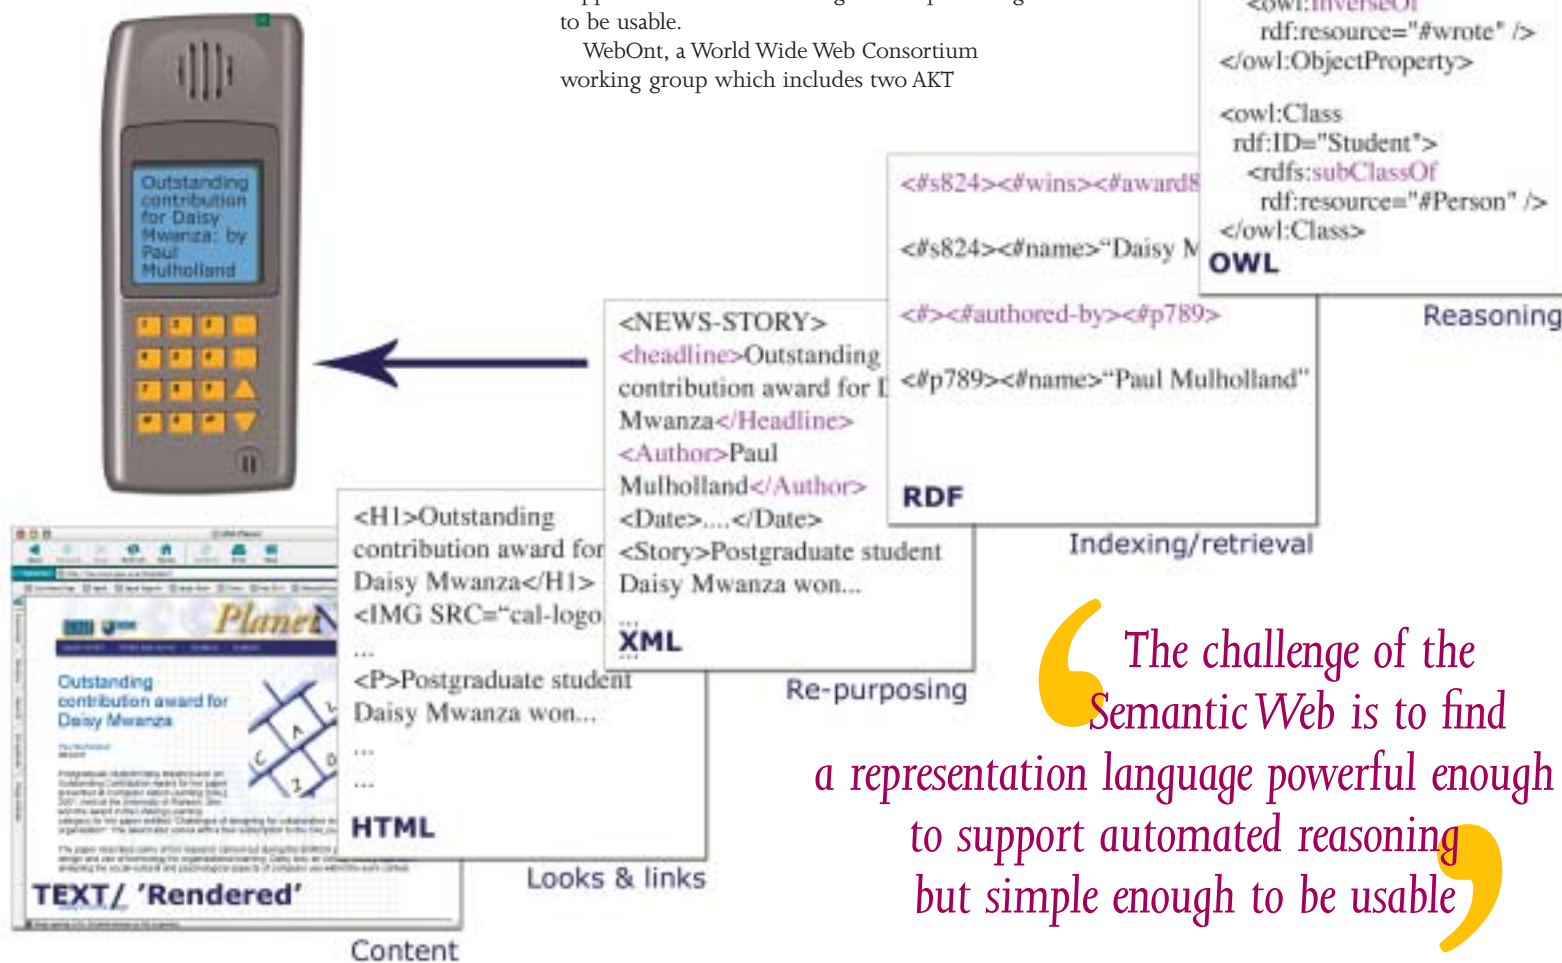
beyond search and retrieval, giving the machine some capacity to understand the implications of particular relations – for instance if an article is authored by a person, then that person wrote the article. It is representation languages like this

that will support the applications described in Re:AKT, such as planning an evening's entertainment in Aberdeen, or recommending computer science papers for researchers to read.

The challenge of the Semantic Web is to find a representation language powerful enough to support automated reasoning but simple enough to be usable.

WebOnt, a World Wide Web Consortium working group which includes two AKT

members, is wrestling with this problem. Only one puzzle remains. If it is so intelligent, why is the Web Ontology Language not called WOL?



Knowledge services

THERE is a huge quantity of information on the web – and our society is getting more and more information-hungry. As the Semantic Web allows increasingly intelligent search and retrieval of information, AKT can begin to provide the technology to satisfy information hunger from the immense resources on the web.

One of the most ambitious AKT projects is to set up an AKTive portal to provide knowledge services for British computer science. There is a huge amount of data about computer science sitting on the various departments' websites, waiting to be Hoovered up – about projects, staff, research interests, courses, prototypes, etc. This data can be used to produce a snapshot of the state of British computer science.

The AKTive portal is an information space organised by an ontology of the important concepts in computing research. Knowledge services which the portal may provide include: recommending relevant people or documents; supporting discussions of published or unpublished papers; issue tracking – the generation and maintenance of a 'to do' list; narrative or template generation for reports or presentations; collecting research intelligence – which projects are on the go in such-and-such a research area.

Once the research is proven, it can be used in fields other than computing.

The information to support all these services is already on the web. AKT's aim is to provide the technology to use this previously unexploited resource.

Amilcare: the key to automating annotation



THE Semantic Web is dependent on the annotation of documents with metadata – all the stuff about the document itself that is so useful for reasoning about it but so hard to collect.

Amilcare, developed by Dr. Fabio Ciravegna at Sheffield, provides the leading solution for

partially automating the annotation process. It is an adaptive Information Extraction engine with the unique ability to handle the whole range of web pages from plain text to structured data.

Into these texts, it can insert annotations indicating the location of significant

information, making it available for subsequent processing.

The software is designed with both novices and experienced professionals in mind, and permits a sophisticated Information Extraction system to be trained and deployed with minimal

input from the user. Amilcare has already found wide acceptance as proved by its integration with a number of Semantic Web tools including CREAM, an annotating and authoring environment from the University of Karlsruhe.

Improving business workflow

MODERN organisations are virtual entities. People are distributed about the place, with different capabilities and responsibilities, all needing to interact and collaborate. Not only that, some of the agents in a business workflow are artificial!

These different agents must get together to achieve complex tasks requiring specialised – read scarce and expensive! – expertise and sophisticated technology.

Co-ordinating resources for optimal efficiency is a hard problem and demands much organisational effort.

Expertise

The AKT consortium has special expertise in this field. The Edinburgh team has a great deal of experience in Workflow and Business Process Modelling and is focusing on the problem of providing an open-architecture

virtual workflow system that is distributed, knowledge-based and agent-based, whose ability can be enhanced by arbitrary intelligent agents where appropriate.

This approach integrates wonderfully with Aberdeen's work on problem-solving agents, and the KRAFT system, a constraint-based problem-solving agent, has been used to integrate with the distributed workflow system.

Other tools across the AKT consortium that might be integrated in this way include, for example, the OU and Southampton's hyperlink tool, process modellers in general, the AKTbus, Netmeeting, Edinburgh's I-X, and Aberdeen's KRAFT system and its inbuilt Constraint-Solver

The ultimate goal of the effort is to provide an environment and means where multiple tools, methodologies and disciplines can be brought to bear to improve organisational performance.

