ERCIM NEWS

European Research Consortium for Informatics and Mathematics
www.ercim.org

Number 53, April 2003

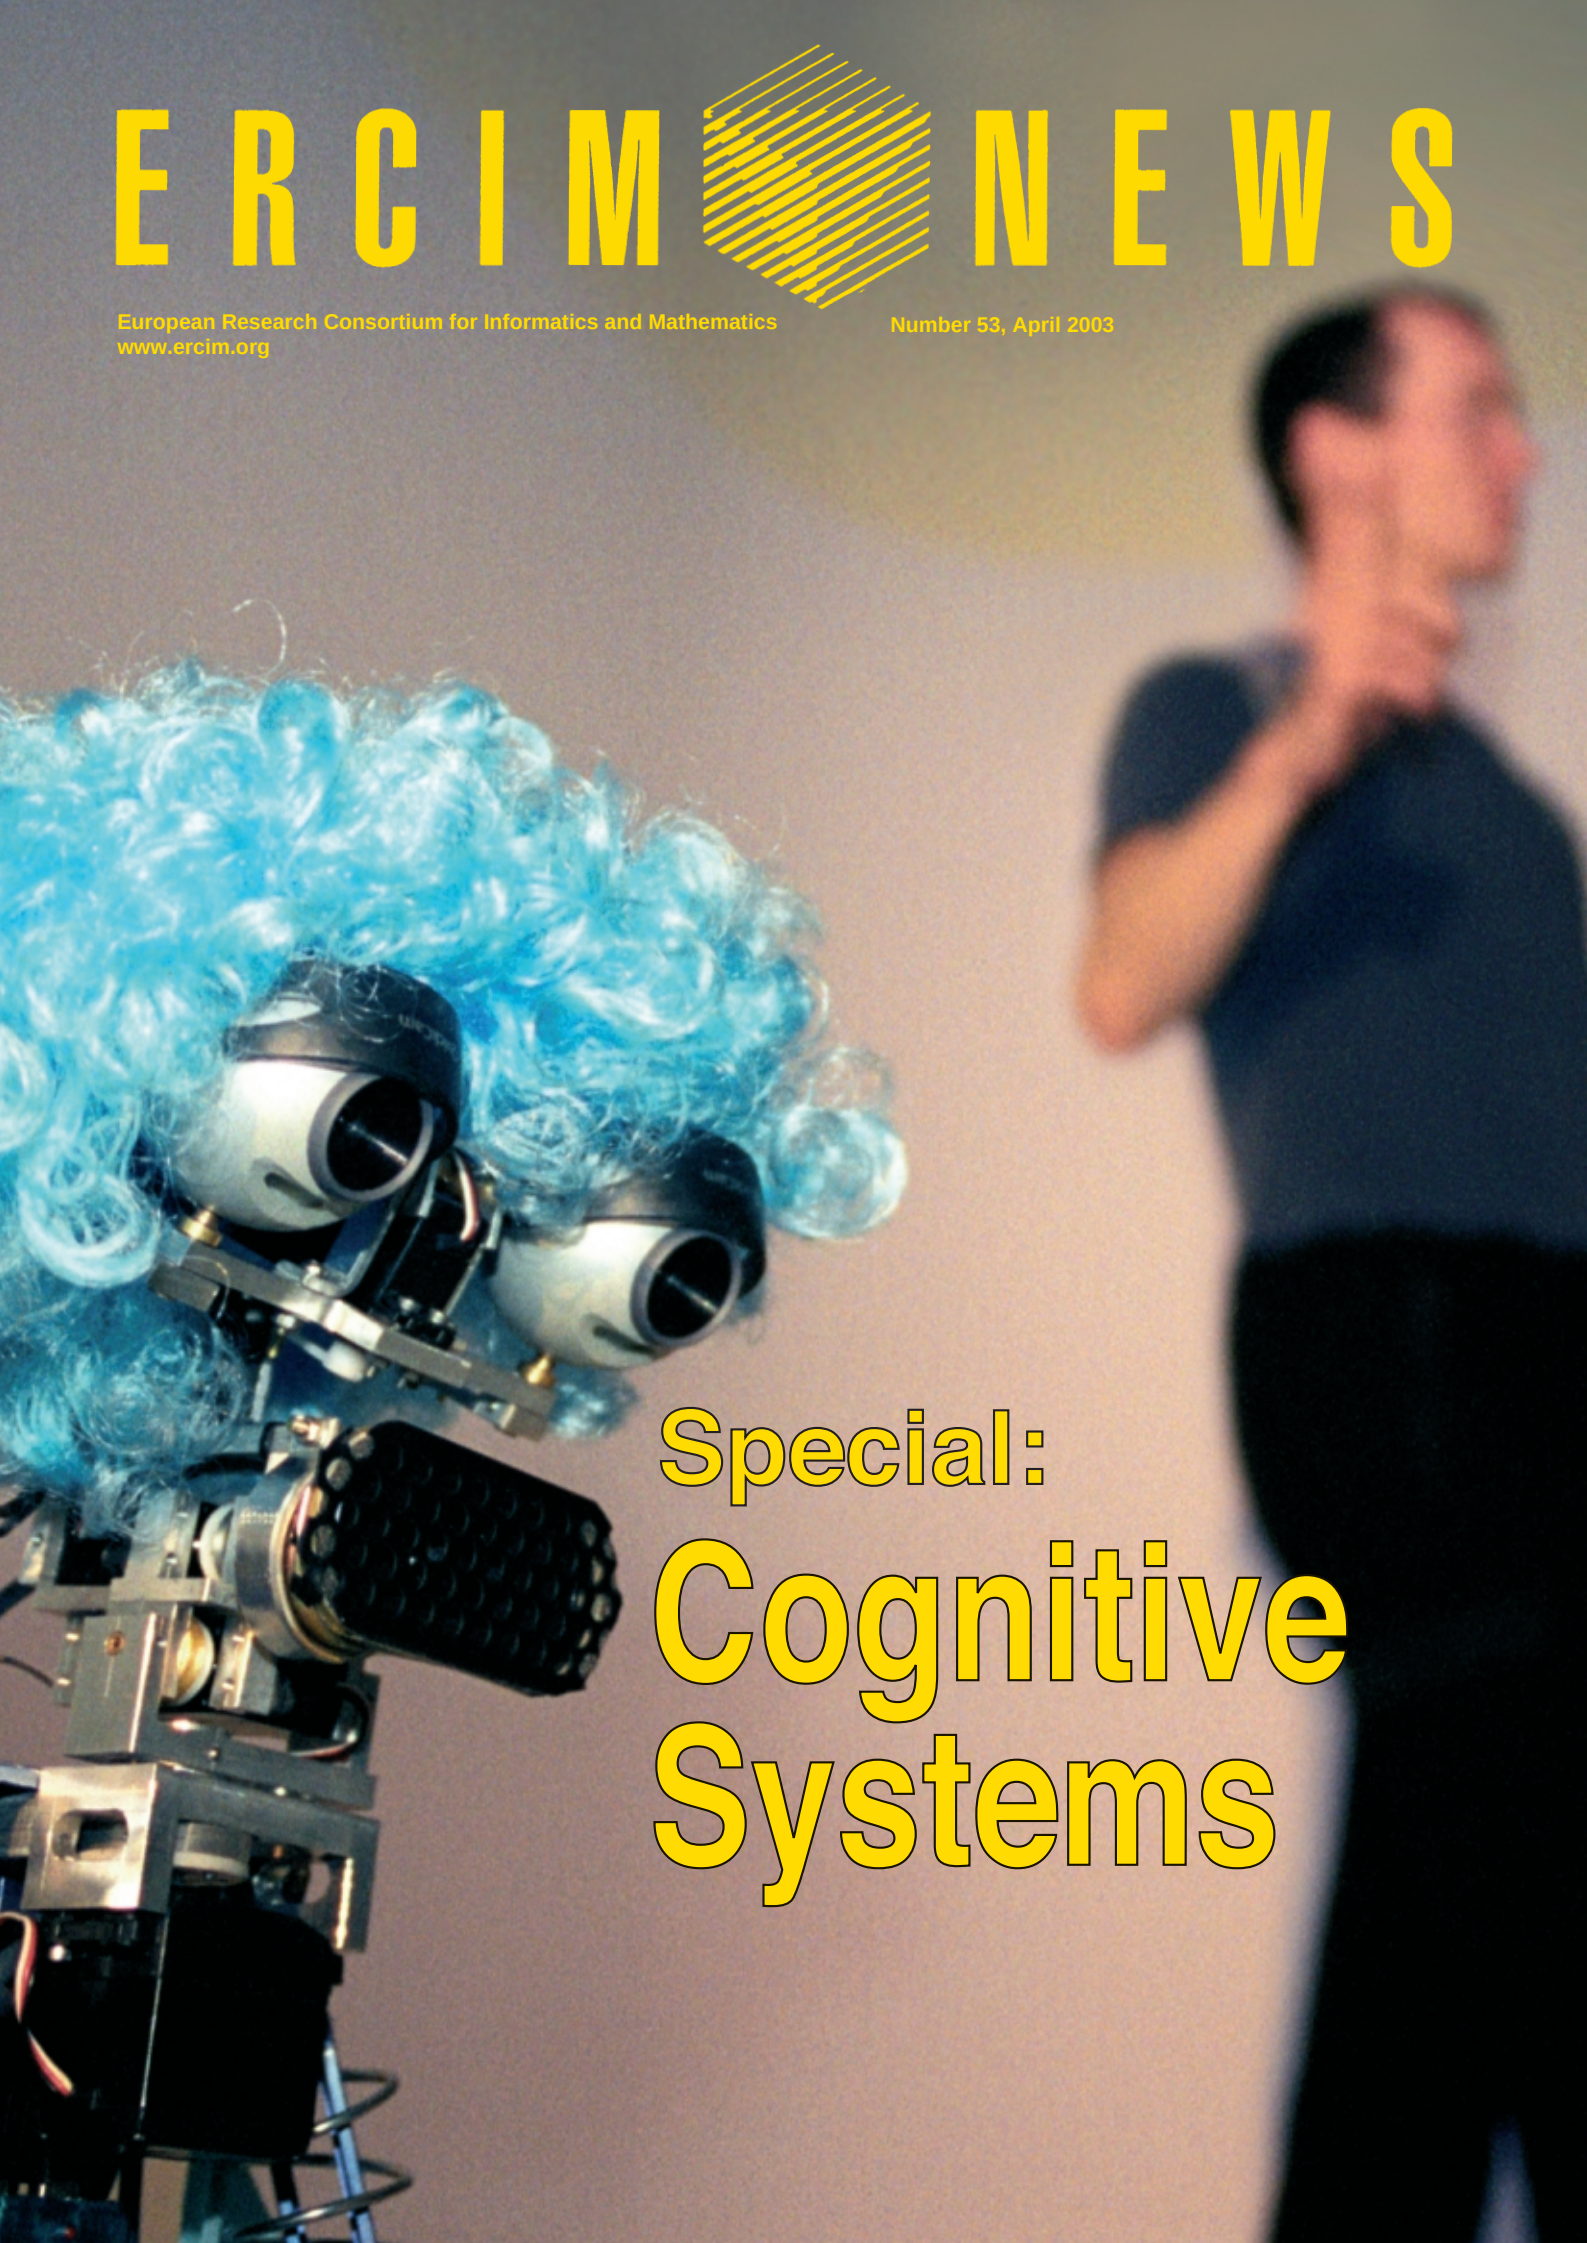

Special:
Cognitive Systems

Contents

Cover: Media Lab Europe's 'JoeRobot',
(see article 'Man and Machine' on page 35).
Photo courtesy of Brent Jones.

Next issue:
July 2003

Special theme:
Applications and Service Platforms for the Mobile User

The Luxembourg economy largely relies on high-level human resources. For the past 15 years, a great number of foreign talent have been attracted by a dynamic situation and business opportunities, in some of Luxembourg traditional fields: mainly financial sector and the associated computer industry, as well as specific technology and communication industries.

To further support the development of technological excellence in Luxembourg, it is essential to build a strong scientific and technological infrastructure that allows cooperation between private sector, higher education and research.

Erna Hennicot-Schoepges, Luxembourg Minister for Culture, Higher Education and Research.

At the Barcelona European Council in 2002, the Heads of State and Government in Europe agreed that research and technological development investment in the European Union must be increased with the aim of approaching 3% of GDP by 2010, a goal to be achieved to a great extent through private sector funding. This goal is very ambitious. In Luxembourg, R&D investment (private and public combined) is currently at 1.8% of GDP - not bad, bearing in mind in this context that public research in Luxembourg is only 15 years old.

In 1999, we established the National Research Fund in order to give a new momentum to research. The National Research Fund was created to help focus R&D efforts in Luxembourg on a number of particularly promising areas and subjects.

These and all the other research programmes and activities in Luxembourg are of course only successful in such a small country because international cooperation has been given the highest priority.

The results of this approach are encouraging: Luxembourg researchers are more successful in gaining European funding, according to the recent European Union's Research Framework Programme. For example, in the Information Society Technology Programme, Luxembourg public research centres' success rate was 33% while European average was 25%.

The National Research Fund joined ERCIM in July 2002. Luxembourg's recent participation in ESA, ESF, EUROHORCS and plans to join EMBC and EMBL are making Luxembourg's research more visible for the international scientific community.

A successful research environment also relies on a good cooperation between higher education and research. A Parliamentary Act on the future development of the University of Luxembourg is at its final stages of preparation and cooperation agreements with Universities and research Centres in Europe and beyond are in preparation in order to establish joint academic curricula.

In this context, let me underline the creation of doctoral programmes of excellence. This initiative aims at meeting several of the challenges I previously referred to :

- to develop high-level doctoral research in a few particular areas, identified as the backbone of current and future Luxembourg economy, among them ICT
- to build a strong partnership between academic research and private sector
- to strengthen international research cooperation
- to propose real platforms for multi-disciplinary international research
- to establish a flexible legal and institutional framework within the future University of Luxembourg.

ITC being particularly critical for the future development of higher education, research and economy, we recently launched a doctoral programme of excellence in this field, called 'LIASIT', Luxembourg International Advanced Studies in Information Technologies.

A survey conducted among Luxembourg companies potentially interested in doctoral research in the field of ITC shows a very high interest in this initiative. LIASIT programme therefore successfully started in September 2002. It currently employs five part-time professors. Nearly ten doctoral projects have been launched since September in close partnership with Luxembourg industries and European universities.

From an organisational point of view, thesis projects are co-financed by public grant and partner contribution. PhD students spend half of their time with the private partner research team and half of their time at LIASIT.

Several types of doctoral projects are emerging :

- 'spin-off linked thesis' focuses on very specific opportunities with high-tech start-up companies involved in EU projects or consortia
- 'industrial thesis' integrates PhD students in classical R&D teams and projects with Luxembourg main industrial and economic actors (eg in such fields as embedded software, signal processing, applied mathematics in finance sector,…)
- 'innovation thesis' enables several companies or public bodies to access collective and shared research. Public transfer centres (Centres de Recherche Publics) play the role of integrators, coordinating these projects.

I am confident that this approach will pave the way for Luxembourg to be a strong partner in the European and international scientific community in the field of computer science and mathematics.

Close Cooperation among Researchers strengthens the Role of Europe in IT

by Michael Krapp

A 'Future Talk' roundtable at the CeBIT fair in Hannover brought together the Directors of the ERCIM Institutes CWI, INRIA and the chairman of the Fraunhofer ICT Group to discuss about ERCIM and the European Research Area.

"The close cooperation among European research institutions is an important advantage compared to the US; Consortia like ERCIM can help to bunch competencies while the leading universities in the US compete one against the other.", said ERCIM manager and INRIA's President Bernard Larrouturou at the computer fair CeBIT in Hannover. In a forum tagged 'Future Talk' and arranged by research institutions along with the German Federal Ministry of Education and Research and the Deutsche Messe AG, Larrouturou took part in a discussion with ERCIM's president Gerard van Oortmerssen and José Luis Encarnação. Encarnação who runs the new Fraunhofer ICT Group in Germany together with his deputy and member of the ERCIM board of directors, Ulrich Trottenberg.

From left: José Luis Encarnação, Bernard Larrouturou and Gerard van Oortmerssen.

Encarnação pointed out that ERCIM could play an important part in the setting up of research consortia. Aiming to bring together the best European researchers in any given field, Encarnação suggested that therefore road maps have to be developed for the advancement of key fields of research in IT. Van Oortmerssen noted that ERCIM has already been organising workshops on behalf of the European Commission and will extend its efforts in providing a leading platform for research cooperation in Europe.

Another topic of discussion involved how to increase young people's interest in studying mathematics and computer science. Van Oortmerssen believes it is necessary to improve the image of computer science and mathematics as disciplines of study. He hopes that the ERCIM fellowship programme will succeed in encouraging young people to study these sciences. Larrouturou's concern is that Europe should put more emphasis in attracting students from other parts of the world. Encarnação reported that in Germany the number of students in information technology is on an upward trend, mostly because of the still excellent career prospects.

The three research managers also discussed the governmental funding of research. All were of the opinion that efforts to translate research results into benefits for the economy must be increased. They also talked about the question of how much basic research will be necessary in information technology. Encarnação described two approaches of doing basic research. One way is to carry out long-term research and look for markets by the time promising results are available. The other way is to look for problems in the economy and to then find solutions through research.

It was also recognized that support for start-up companies is a common goal of research institutions in France, the Netherlands and Germany. Although the time of the gold digger in information technology is over Encarnação, Larrouturou and van Oortmerssen see excellent chances for business with the development of exciting ideas.

ERCIM PhD Fellowship Programme

ERCIM offers 18-month fellowships in leading European information technology research centres. Fellowships are available for PhD-holders from all over the world. Next Deadline for application: 30 April 2003

Fellowships are of 18 months duration, generally spent in two ERCIM member institutes. The fellow will receive a monthly allowance which may vary depending on the country. In order to encourage the mobility, a member institute will not be eligible to host a candidate of the same nationality.

Topics include but are not restricted to the scientific fields covered by the ERCIM Working Groups:

- Applications of Numerical Mathematics in Science
- Constraints Technology and Applications
- Control and System Theory
- Dependable Software-Intensive Systems
- Digital Libraries
- E-Learning
- Environmental Modelling
- Formal Methods
- Health and Information Technology
- Image and Video Understanding
- Matrix Computations and Statistics
- Soft Computing
- User Interfaces for All.

A detailed list of topics in which ERCIM institutes are active is available at http://www.ercim.org/activity/expertise.html

Conditions

Applicants must:

- have a PhD degree or be in the last year of the thesis work
- be fluent in English
- be discharged or get deferment from military service.
- start the grant before January 2004
- have completed their PhD before starting the grant.

How to apply

Applications must be submitted online
A detailed description and the application form are available at:
http://www.ercim.org/fellowship/

Euro-Legal

News about legal information relating to Information Technology from European directives, and pan-European legal requirements and regulations.

e-Copyright

The recent case of Sony Computer Entertainment v Paul Owen & Ors appears to have established the right of copyright owners to use copy protection technology in their software, and manufacturers/distributors of devices designed to circumvent this protection can be sued for infringement, even if those devices have other lawful purposes.

Sony embedded a code into its PlayStation 2 games which restricted read access through SONY consoles. Additionally the code was customised according to the regional market of sale, ie, Japan, US or Europe. The defendant imported a chip device known as "Messiah" which when inserted into the PlayStation 2 could circumvent the copy protection and by-pass these regional codes. The defendant argued that the device could be used for a number of lawful purposes, such as providing back-up copies in the event the original CD became damaged or destroyed. The Judge disagreed and held that copyright is territorial and that games exclusively licensed to a particular country cannot lawfully be imported for use in another without the copyright holder's permission. The license of a Sony games holder covered playing the game on a particular disk and console and was not for the purpose of making back-up copies for normal use.

This ruling is in contrast to the general rule established by the House of Lords in 1988 which held that selling high speed tape copying equipment did not breach copyright, since the equipment could be used to copy tapes in many other ways.

Defamatory Statements

An Australian High Court has ruled that the Dow Jones can be sued in Australia for an article about an individual which it published on a server in New Jersey, USA, on the basis that the individual's reputation in Australia had been damaged. This could open up a way for any international new agency to be sued in the Australian courts, even if the individual is not resident in Australia. All they have to prove is that they have a reputation there to defend. If this ruling is copied in other countries, it could open a flood-gate of potential litigation for on-line publishers who could be sued in every country in which a defamatory article is read on-line if they are unable to prove that the article is legal or could be justified. Publishers would have to adopt the standard required by each jurisdiction in which their on-line material could be accessed, even though there would be differing standards of evidence and different defenses. If this were to extend to newspaper publishers it may result in a withdrawal of publications rather than facing the potential of liability everywhere the web can be accessed.

Domain Name Dispute

WIPO has ruled that domain names registered in "bad faith" must be given up. Robbie Williams brought the case against Howard Taylor, who registered the domain name RobbieWilliams.info but pointed it to the website of the rival band Oasis. The Dispute Resolution Panel found that the domain was registered in "bad faith" on the basis that it was not the "action of a fan, but rather the action of someone who wished to provoke the Complainant". The judgement reinforced the rights of celebrities and businesses to their intellectual property and is a warning against the malicious online use of trademarks in an infringing act.

by Heather Weaver,
CCLRC, Tel: +44 1 235 446151, E-mail: H.Weaver@rl.ac.uk

News from W3C

XML Fifth Birthday

The Extensible Markup Language (XML) celebrated the fifth birthday on 10 February 2003. XML was first published as a W3C Recommendation on 10 February 1998. Since its first introduction, the Extensible Markup Language has become pervasive nearly everywhere that information is managed. With its companion and follow-on specifications, XML has changed not only the way people publish documents on the Web but also the way people manage information internal to their enterprise. Dave Hollander and C. M. Sperberg-McQueen, participants in the W3C XML Working Group who wrote the original twenty- five page XML specification, write about XML's growth in an article at http://www.w3.org/ 2003/02/xml-at-5.html. The authors believe, "Just as interchangeable parts drove the Industrial Age, reusable information powers the Information Age."

Liam Quin, W3C XML Activity Lead noted:
The earliest customers of the XML standardization process were, for the most part, technical writers: people working on using and publishing technical documentation and doing large-scale document engineering. They wanted to deliver richer content to their customers, and they wanted to do it on the World Wide Web.

The move to richer markup for documents has been slow in coming, but XML and XHTML support in most modern web browsers is helping to increase momentum.

What many of us hoped for, but perhaps did not dare to expect, was that XML would take off as a general-purpose format for structured textual information. In fact, XML has become a part of general computing infrastructure.

Applications of XML range from configuration files through to remote procedure calls, from desktop menu definitions to chat protocols. With every new usage of XML, the value of all the interoperable XML tools, both proprietary and open source, is increased. As the value of the tools increases, so does the value of XML itself.

Five years ago, if you suggested to a programmer that a configuration file be stored in SGML, you'd probably have got either a blank stare or a hostile reaction. Today, XML is used by the GNOME desktop on Linux, by the Jabber chat protocol, and is even central to the upcoming release of Microsoft Office. We've made it. We're mainstream.

Of particular interest on the World Wide Web is the success of XSLT, the XML way of transforming documents. This XML-based language has proved useful as a glue, as middleware, connecting among other things databases, text documents, web servers, web browsers and styles. It's also finding a place as part of Web services, helping to process business transactions.

Two early design decisions with XML were to minimize the number of optional features, and to require strict error checking. The result of these decisions has been that there is a very high degree of interoperability between tools. Another decision was to use Unicode as the document character set: work on accessibility and internationalization are helping to make the World Wide Web truly a place for all people.

In the past five years XML has gone from an obscure format for technical manuals to a central part not just of the World Wide Web but of modern computing and business. The World Wide Web Consortium continues to provide a leading role in guiding this free and universal technology towards maturity.

Last W3C Recommendations

- Scalable Vector Graphics (SVG) 1.1 Specification
- Mobile SVG Profiles: SVG Tiny and SVG Basic
- Document Object Model (DOM) Level 2 HTML Specification

An exhaustive list of W3C Recommendations is available at: http://www.w3.org/TR/#Recommendations

Hadden Award for Judy Brewer

Judy Brewer received the Susan G. Hadden Pioneer Award 'for pioneering efforts in telecommunications and consumer access' from the Alliance for Public Technology (APT) in Washington, DC, USA on 21 February 2003. The award recognizes those who continue Hadden's legacy of ensuring equitable access to technology as a democratizing principle. Judy Brewer is Director of W3C's Web Accessibility Initiative (WAI). For more information about the WAI, see http://www.w3.org/WAI/

W3C Days in Beijing

A two-day presentation of W3C will be held in Beijing 16-17 April 2003. Attendees of the China International Forum on WWW's Development can participate in panels on accessibility, SVG, mobile Web, the Semantic Web, VoiceXML and internationalization. The event is co-organized by the China Computer Federation and the W3C Office in Hong Kong. http://www.w3c.org.hk/CIFWeb03/

WWW2003

The International World Wide Web Conference, this year to be held in Budapest, Hungary (see announcement on page 54), provides a public forum for the W3C through the annual W3C Track. Speakers from W3C, including Tim Berners-Lee, present W3C Activities and recent developments. http://www2003.org/t_www.htm

ERCIM is the European host of W3C.
For more information about W3C, see http://www.w3.org/

Extending Digital Libraries towards Multi-Disciplinary Communities

by Tom Gross and Wido Wirsam

The CYCLADES open collaborative virtual archive environment is a virtual community system that supports scholars in searching and managing e-print documents from a large number of digital libraries. Users can share their results in existing communities as well as find new collaborators with similar interests.

The CYCLADES system provides an open collaborative e-print archive environment to support online communities of scholars. Electronic pre-print (e-print) archives are important vehicles for the dissemination of preliminary results and non-peer-reviewed grey literature. Most of them focus on the dissemination of information within disciplinary or institutional communities. Scientific research, however, tends increasingly to be interdisciplinary. There is thus a growing need for easy retrieval of information from diverse sources, and for communication and collaboration across traditional community boundaries. CYCLADES addresses these issues and supports the transition of e-print systems into the genuine building blocks of a transformed scholarly communication model by developing a set of leading-edge technologies.

The CYCLADES system integrates a set of functionalities that support the user when accessing very large virtual e-print archives with:

- functionality for efficient and effective retrieval of relevant information from many large, distributed and multi-disciplinary digital archives
- feedback on the degree of relevance of the retrieved information
- regular information about new publications in the archive environment which are relevant to the users' interests
- automatic retrieval of users' long-term information needs (user profiling)
- rapid dissemination of the search results world-wide.

A special set of features provides communities of scholars with functionality for:

- dissemination of relevant information to community members in the form of recommendations, which are based on collective profiles and behaviour
- quick on-line annotations on research results published by members of the community
- carrying out community services such as peer review, which requires the annotation of on-line manuscripts by reviewers, and the sharing of these annotations among editors, authors, and others
- enabling community members to learn from, contribute to, and collectively build upon, the community's knowledge.

The archive environment is composed of a large number of heterogeneous, multi-disciplinary archives and supports interoperability between them. It is based on the Open Archive initiative (OAi) (cf. http://www.openarchives.org/). The OAi consists of a technical and organisational framework designed to facilitate the discovery of content stored in distributed e-print archives.

The CYCLADES system consists of modular software packages, which are executed on geographically distributed computer systems. The different services communicate over an interface that is based on modern standardised internet communication technology (XML-RPC). This architecture allows the different servers to be optimised to achieve best performance in the specific task for which it is responsible. Highly time consuming operations (eg re-indexation of archives) can be performed in the background without decreasing the overall system performance. This specification also gives excellent flexibility in the adaptation, updating or optimisation of services or the integration of a new functionality.

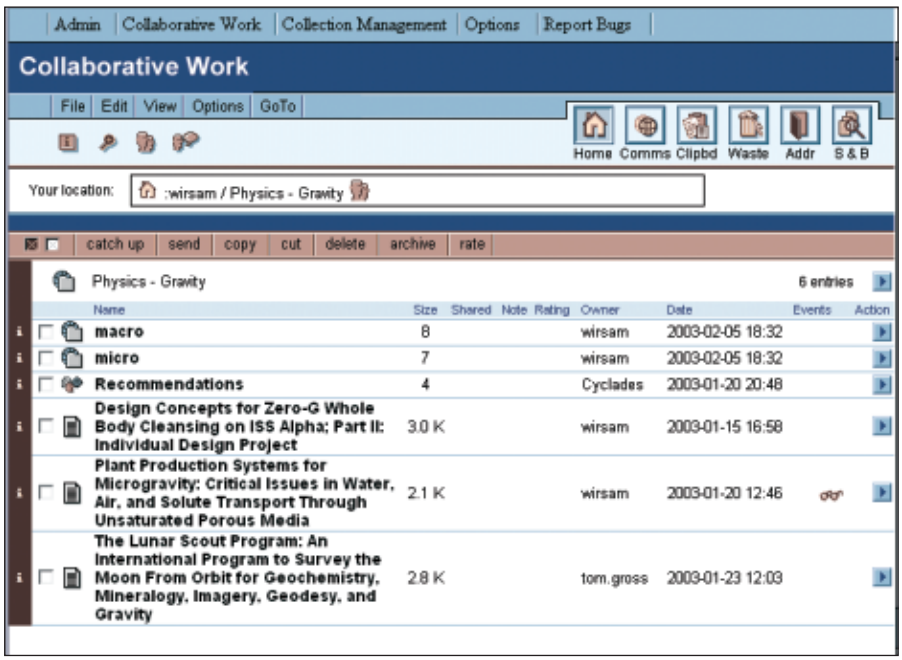

The CYCLADES collaborative working environment.

The service responsible for accessing all connected archives is the Access Service (AS). It harvests, indexes and retrieves metadata records from the underlying archives.

The *Collaborative Work Service* (CWS) provides a folder-based environment for managing metadata records, queries, collections, external documents, and annotations. The CWS supports users with functionality for the management of folders and their contents and for cooperation among community members, including special awareness features reflecting the groups' activities in the shared information space.

The *Rating Management Service* (RMS) is an internal component of the CWS that stores ratings and provides a query interface to other services of CYCLADES.

The *Search and Browse Service* (SBS) supports searching for records from the various collections, formulating queries, and browsing attribute values and metadata records.

The *Filtering and Recommendation Service* (FRS) provides personalised filtering of queries and query results, provides recommendations of relevant records, collections, users and communities. The FRS analyses a user's behaviour and data, generates a user profile describing the user's interests, and provides information according to this profile.

The *Collection Service* (CS) supports the creating and editing of collections. Collections partition the information space according to the users' interests, and make the individual archives transparent to the user. The CS allows each user to create and maintain a personal list of collections consisting of a sub-set of all the collections available.

The *Mediator Service* (MS) integrates the other services of CYCLADES and acts as a registry for them. Each service registers with the MS. Whenever a service needs to communicate with another service, it asks the Mediator Service for a list of services of the appropriate type. The MS also manages the registration and login of users to the CYCLADES system.

The CYCLADES project, partly funded by the European Union (IST-2000-25456), is administrated by ERCIM. The work is carried out by the Fraunhofer Institute for Applied Information Technology (FIT), CNR, FORTH and the University of Duisburg. The project started in February 2001 and will run until July 2003. At this stage, the implementation, integration and testing of the different services forming the environment has been completed and the user evaluation is starting.

If you are interested in CYCLADES and would like to try the system or participate in the evaluation, please follow the link below or contact the authors.

Links:
http://www.fit.fraunhofer.de/cyclades
http://www.ercim.org/cyclades/

Please contact:
Tom Gross, FhG-FIT
E-mail: tom.gross@fit.fraunhofer.de

Umberto Straccia, ISTI-CNR
E-mail: straccia@iei.pi.cnr.it

ERCIM-sponsored Events

ERCIM continues sponsoring up to twelve conferences, workshops and summer schools per year. The funding for all types of events is in the order of 2000 Euro.

Conferences
ERCIM invites sponsorship proposals from established conferences with an international reputation, where substantive overlap can be shown between the conference topic and ERCIM areas of activity. Typical cases would include annual conferences in computer science with international programme committees, substantial international participation, and proceedings published with an established international science publisher.

Workshops and Summer Schools
ERCIM sponsors workshops or summer schools (co-) organised by an ERCIM institute. The additional funding provided by ERCIM should be used to enhance the workshop by, for example, increasing the number of external speakers supported.

Forthcoming events Sponsored by ERCIM:
- WWW2003 - The Twelfth International World Wide Web Conference, Budapest, 22-24 May 2003 (see announcement on page 58)
- 2nd International Conference on Universal Access in Human - Computer Interaction (UAHCI 2003), Heraklion, Crete, Greece, 22-27 June 2003
- European Conference on Object-Oriented Programming (ECOOP), Darmstadt, Germany, 21-25 July 2003
- ECDL2003 - 7th European Conference on Research and Advanced Technology for Digital Libraries, Trondheim, Norway, 17-22 August 2003
- FM 2003: the 12th International Formal Methods Europe Symposium, Pisa, Italy, 8-14 September 2003
- CP 2003 - Ninth International Conference on Principles and Practice of Constraint Programming, Kinsale, Ireland, 29 September - 3 October 2003.

For detailed information about ERCIM event sponsorship, see:
http://www.ercim.org/activity/sponsored.html

RESET - A Roadmap for European Research in Smartcard Technologies

by Olivier Trebucq and Bruno Cucinelli

RESET is the first ever made attempt of the smart card industry to assess, in a global way, its technology priorities and R&D orientations. Experts in six technology domains have organised workshops from October to December 2002 and are now finalising the roadmap.

RESET is a Thematic Network to build a roadmap investigating future research challenges and opportunities in the field of smart cards. It includes representatives from European smart-card industry and the research community. RESET is jointly managed by ERCIM and Eurosmart, supported by the IST Programme of the European Commission.

Smart cards are key components addressing security needs in a number of well established consumer applications. Although the smart card industry is currently experiencing a decreased development rate - mainly resulting from the global slow down of ICT markets and in particular mobile phone markets - experts are convinced that there is still an enormous potential for smart card deployment in traditional, and new application areas.

Traditional high-volume applications include banking, telecom, pay-TV, etc. There is still a strong need for innovation to address the requirements of those applications to overcome existing limitations and anticipated evolving environments. But an even higher potential for smart card technology applications is expected from upcoming ubiquitous computing and ambient intelligence environments. These environments create a need for 'handy' personal keys that can provide the required level of trust and confidence to users in networked applications. Current smart card technology has an advance over other solutions to address these needs - however some key aspects such as improvement of online security management and speed of communication protocols must be seriously considered to enable the integration of a trusted personal device such as a smart card in networked applications.

Domains of Investigation In order to establish a suitable roadmap, the project had to take into account a specificity of smart card systems, which is to integrate a wide range of technologies. The investigation carried out has therefore been divided in six main technology areas, each of them covered by one expert working group:

- communication and networking
- systems and software
- interface technologies
- peripherals, subsystems and microsystems
- high-end cryptography, tamper-resistant and security technologies
- micro-electronics.

The investigation of the working groups was carried out following a common framework which included the following topics:

- state of the art: existing and emerging technologies, their limitations, competing technologies
- on-going research: inside and outside the smart card world
- evaluation of technology and marketing requirements
- research orientations for improvement: short, medium and long term.

The results of this investigation and consultation process will be published in a report. A draft has been presented and discussed during an public seminar in Brussels on 3 April. The report starts with a short description of the socio-economic context that influences the development and deployment of new generation smart card applications. Then a summary of the main outcomes of the 6 working groups is introduced. The report also includes a summary of the main driving and blocking factors that condition the evolution of smart cards and then identifies the technical challenges that need to be mastered to enable the European smart card technology suppliers to deal with these factors, address anticipated requirements and exploit new market opportunities. Further, it lists the resulting main R&D targets and RTD programmes and resources that have a capacity to carry out smart card technology related research. Recommendations are proposed for integrated collaborative research programme that would federate the required level of European R&D resources to address these R&D targets.

The final report will be available on the project website. Interested people can submit their email address to be informed when the report is published.

Link:
http://www.ercim.org/reset
http://http://www.eurosmart.com/

Please contact:
Bruno Le Dantec, ERCIM office
E-mail: office@ercim.org

Cognitive Systems: Towards an Integration of Symbolic and Sensor-Motor Intelligence?

by Cristiano Castelfranchi

The European Commission has identified Cognitive Systems as one of the priorities for the new generation of research projects to be developed from 2003 to 2008 (http://www.cordis.lu/ist/workprogramme/fp6_workprogramme.htm). The stated objective is to construct physically instantiated or embodied systems that can perceive, understand (the semantics of information conveyed through their perceptual input) and interact with their environment, and evolve in order to achieve human-like performance in activities requiring context-(situation and task) specific knowledge.

ERCIM News has chosen to devote a special issue to this exciting research challenge in order to monitor what is under development in Europe (but not only in Europe), and what is the current status of research and development in this domain.

Over the last twenty years, in the field of Cognitive Science, there has been a real scission between the symbolic paradigm, on the one side, and the heterogeneous army of its enemies: neural nets and connectionism, embodied and situated intelligence, symbol grounding, dynamic and evolutionary approaches, distributed cognition, on the other. In the meantime, important changes have taken place. Taking Artificial Intelligence as an example, we can cite the reconciliation between reactivity and planning with the BDI approach (ie, modelling the mind and its processing in terms of Beliefs, Desires, Intentions), or between a reasoning mind and a open uncertain world; the idea of a decentralised and emerging intelligence among distributed interacting systems; the development of new formalisms for autonomous action and cognitive robotics; or the introduction of so-called 'affective computing'. At the same time, more radical and simplistic attacks, such as the elimination of planning, or the elimination of intelligence as manipulation of internal symbolic representations, or the rejection of any modularity of mind, have been relaxed. Thus, a less ideological confrontation, and perhaps even a reconciliation and integration is now possible. It is clear that Cognitive Systems, able to deal with a physical, embodied and situated environment, but also endowed with mental representations, able to use language, and to exploit both experience (learning) and inference (reasoning), reactions and planning, will be a privileged area for such an attempt.

Indeed, Cognitive Systems is a really important research area from both scientific and technological perspectives, dealing with physical environments and mental representations, and exploiting both experience (learning) and inference (reasoning) in order to be able to autonomously explore, learn, react, decide, solve problems, coordinate with each other, interact with humans, etc. There are numerous and important applications under way, from space exploration to emergency handling, from industry to domestic assistance, to worrying military applications. However, all existing projects and working systems are partial and incomplete when compared with the ambitious model targeted in the Call of the Sixth Framework Programme. The aim of this special number of ERCIM News was thus to provide a broad and sufficiently representative view of current advanced attempts to build embodied Cognitive Systems.

Systems can be incomplete in different ways and there are two partially independent paths towards completeness aiming alternatively at:

a) obtaining a working, integrated, complete picture of an internal cognitive architecture, and of its components and functions (ie decision and reaction, inference and learning, emotions and 'rational' reasoning);
b) completing the coupling between the agent and its environment in order to have an effective, complete loop: perception, cognitive processing, (re)action, effects or independent changes in the word, perception of changes, and so on.

Which of those paths is more important, or should be given precedence? We believe that they should advance in parallel - as is in fact

happening - and preferably interact as, on the one hand, a more effective coupling with a complex environment may well require more complicated cognitive architectures and richer cognitive functions, while, on the other - cognitive components and architectures must be functionally justified by showing their advantage and adaptive value when interacting with a constraining environment.

This perception also clearly emerges from the contributions in this issue, which range from neural nets to logic for action, from reactivity to planning, from situated to distributed approaches; and, in some cases, explicitly discuss (although briefly) the issue of anti-representationalism. Two very interesting and converging movements can be witnessed: on the one hand, some research groups are going from sensor-motor, procedural, and situated levels of cognition towards 'more abstract representations'; on the other, there is a trend moving from merely symbolic reasoning towards more grounded, embodied and sensor-motor representations. We can observe a remarkable and explicitly interdisciplinary effort, involving psychology, neuroscience, and philosophy among others.

Thus the picture of Cognitive Systems - as emerges here - not only has various focuses (mainly: mind-body integration in cognitive architecture; multi-agent environmental co-ordination) but clearly shows the fragmentary and incomplete status of this domain. As yet, we have no complete, mature or fully integrated architecture or model; we continue to have several 'isolated components'; exploration for new models is enthusiastic but sometime a bit naive; there is a strong, useful, but premature desire to model technologies for interesting applications. In a word, this is an exciting, promising, growing domain, that will become mature during and (also) thanks to European initiative in the Sixth Framework Programme.

Please contact:
Cristiano Castelfranchi, ISTC-CNR
E-mail: castelfranc@unisi.it

ARTICLES IN THIS SECTION

The Many Faces of Behavior

by Torbjørn Dahl

At the University of Southern California (USC), the Interaction Lab under Professor Maja Mataric takes a multi-pronged approach to studying embodied systems and the elusive concept of 'behavior'.

As a successful and growing part of USC's Center for Robotics and Embedded Systems, the second largest center for robotics in the world, the Interaction Lab has the resources to attack the problem of embodied behavior on many fronts. The research program includes the recognition, representation, adaptation, transfer, and synthesis of behavior. The application domains span a wide range of embodied systems and use hardware as diverse as humanoids, robot dogs, and groups of mobile robots. Anchored in the Behavior-Based (BB) control paradigm, the Interaction Lab's work is developing BB control into a mature field rich in applications and connections to other frameworks for analysis and control.

Recognition, Adaptation, and Synthesis of behavior place different requirements on the internal representation of behavior within a robot. Here we present results from three current projects. The first project, Primitive-Based Imitation Learning, studies the role of primitives, including oscillatory, discrete, and postural primitives, in Imitation Learning. The second project, Natural Methods for Human-Robot Interaction, demonstrates imitation learning and the transfer of behavior between humans and robots. The final project, Scheduling with Group Dynamics, demonstrates a distributed task-allocation algorithm that allows a group of autonomous robots to optimize their performance in environments ill-suited for classic scheduling algorithms.

Primitives-Based Imitation Learning

Rather than explicitly planning the trajectory of motion of a limb anew for each movement, this project is inspired by theories that posit that human beings tend to chose from a limited but possibly large repertoire of movement primitives. The inspiration for primitives-based motor control is based on evidence from experiments carried out on frogs and rats, where complete movements (such as reaching and wiping) could be produced by potentiating an electrode in different regions of the spine of spinalized frogs.

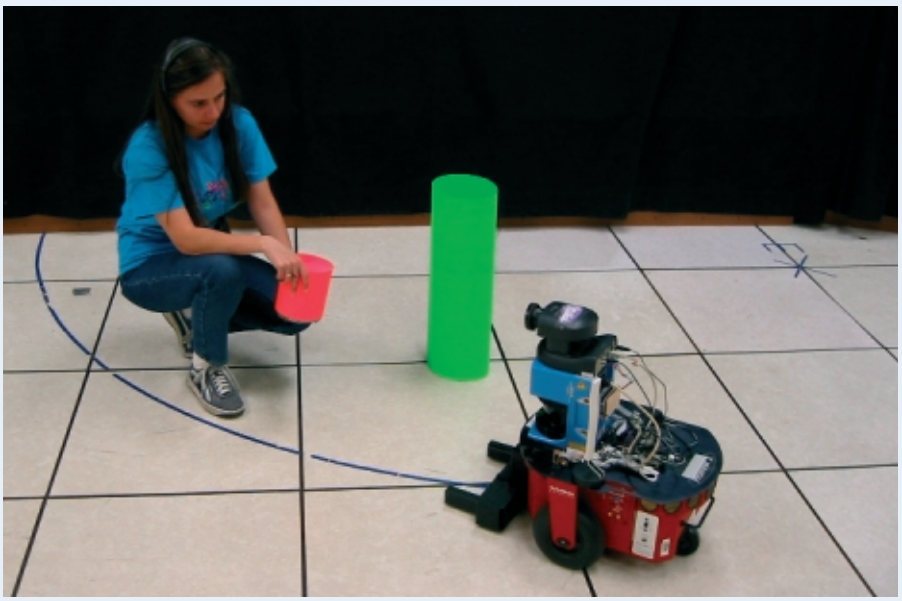

© USC Robotics Lab.

A Pioneer mobile robot learning from a human teacher.

Primitive-Based Imitation Learning, focuses on developing a model of learning by imitation. As one of the most powerful yet poorly understood forms of learning in nature, imitation presents an important research problem in AI and machine learning, as well as in the behavior and neural sciences. The recent successes of this project include new technologies for motion capture and algorithms for extracting primitives automatically. Currently active members of this project are PhD students Evan Drumwright, Chad Jenkins, Amit Ramesh, and Monica Nicolescu.

Natural Methods for Human-Robot Interaction

The goal of this project, mainly the work of PhD student Monica Nicolescu, is to extend a robot's model of interaction with humans so that it can induce changes in a human's behavior and also express its intentions in a way that humans can easily understand. The challenge the project addresses is allowing interaction to occur without the need for an explicitly shared vocabulary between the robot and the human. One important advantage of using a body language technique is that the method is not restricted to having a robot with a humanoid body or face: the approach does not require structural body similarities between the interacting agents in order to achieve successful interaction. Even if there is no exact mapping between a mobile robot's physical characteristics and those of a human user, the robot is still able to convey a message to the human. The approach relies on implicit interaction, which can be achieved by designing an applicable subset of body movements and behaviors for the robot which are already known and understandable by humans through their common sense.

Multi-Robot Task Allocation through Vacancy Chains

This project tackles the problem of group dynamics in Scheduling and is mainly the work of Torbjorn Dahl, a post-doctoral research associate. Existing task allocation and scheduling algorithms, including task-allocation algorithms for multi-robot systems, generally assume that tasks are independent. This assumption is often violated in groups of cooperative mobile robots, where the group dynamics can have a critical impact on performance. The project has developed a multi-robot task allocation algorithm that is sensitive to group dynamics. An algorithm based on vacancy chains, a resource distribution process common in human and animal societies, has been developed.

In particular, the project has studied the problem of cooperative transportation. It has been demonstrated, through experiments in simulation, that if robots keep local task utility estimates, and follow a greedy task selection policy, the interactions in the group cause the collection of learned policies to converge toward an optimal allocation pattern as defined by the vacancy chain framework.

As the robots are continuously updating their individual utility estimates, the vacancy chain algorithm has the additional property of adapting automatically to changes in the environment, eg, robot breakdowns or changes in task values. Experiments show that in the case of such changes, the vacancy chain algorithm consistently outperforms random and static task allocation algorithms.

Links:

USC the Interaction Lab: http://robotics.usc.edu/~agents

USC's Center for Robotics and Embedded Systems: http://cres.usc.edu

Primitives-Based Imitation Learning: http://robotics.usc.edu/~agents/Links/Research/imitation.html

Natural Methods for Human-Robot Interaction: http://robotics.usc.edu/~monica/Research/HRI/hri.html

Multi-Robot Task Allocation through V acancy Chains: http://robotics.usc.edu/~tdahl/projects/vacancychains.html

Please contact:
Torbjørn Dahl, Interaction Lab,
University Southern California, USA
E-mail: Torbjorn-Semb.Dahl@ffi.no

IDA, a Software Agent Cognitive System

by Stan Franklin

A US Navy sailor writes to IDA, "I'm approaching the end of this tour of duty and will need a new job. Can you find me something suitable in the San Diego area?" A couple of months of negotiation via a dozen or so email messages follow. "Can I have the job at the El Toro Marine Base?" "No, you must go to sea duty, and that temporary detached duty won't count." "Can I wait two weeks for the next requisition [jobs] list." "Sure, maybe something more suitable will come up." Issues concerning the sailor's preferences, job qualifications, and Navy policy are discussed. Finally the sailor agrees to an offered job aboard the USS Pennington based out of San Diego. A negotiation between a sailor and a human personnel officer? No, between a sailor and IDA, a software agent cognitive system.

IDA's environment consists of a typical fast workstation with lots of memory, a connection to the Internet, and access to several US Navy databases for personnel records, training schedules, etc. She senses her environment by interpreting incoming strings of symbols from email messages and database records, and acts on it by composing and sending her own such email messages and database queries. She pursues her own agenda, all of which has to do with persuading a sailor to accept a job she has selected as suitable and offered. IDA is a software agent that goes a long way towards being a cognitive system. Her job is to completely replace a human 'detailer,' the personnel officer who normally performs such negotiations.

Supported by the US Navy over the past five years, the IDA project is the work of the 'Conscious' Software Research Group at the University of Memphis. It pursues two objectives: 1) to automate the Navy's assignment of personnel, and 2) to model human cognition in the process. 'Conscious' in the sense that she models a psychological theory of consciousness, Baars' global workspace theory, IDA also models theories of perception, memory, decision making, deliberation and constraint satisfaction. She is a fertile source of hypotheses about human cognition that can be tested by cognitive scientists and neuroscientists.

The same IDA technology is capable of automating the tasks of many human information agents, that is people who negotiate in natural language, access information in databases, adhere to company or agency policy, make decisions, and create paper products. Such human information agents might include customer services agents, travel agents, loan officers for a bank, and various low

to mid-level public servants. Though the IDA technology would be common to all of these, a considerable chore of gathering and encoding appropriate knowledge would be required to implement any one of these.

This IDA technology is based on a host of disparate mechanisms taken from the 'new' artificial intelligence. These include the Hofstadter and Mitchell's Copycat architecture, Kanerva's sparse distributed memory, Maes' behavior nets, and Jackson's pandemonium theory. IDA is currently up and running, and has been tested to the satisfaction of Navy detailers. Watching IDA in action, their reaction is typically a nod of the head together with "Yes, that's how I do it."

This IDA technology is currently being adapted to produce a web-based multi-agent system in which each of the more than 300,000 sailors will be represented by his or her own IDA-like agent devoted to the sailor's interest. The sailor will communicate with his or her agent in a private chat room environment on the web. Each of the Navy's commands will also be represented by an IDA-like agent looking out for its needs. Between periodic consultations with its sailor, a sailor agent will negotiate with various command agents until some agreement is reached. If both the sailor and the command approve the agreement made by their agents, it's taken before an IDA-like broker agent who is concerned with adherence to Navy personnel policies. The broker agent's approval constitutes the final blessing on the agreement. The sailor will then receive orders to the agreed upon job.

If a multi-agent system consisting of hundreds of thousands of quite sophisticated software agents seems impossible, that's because it is impossible. To make this project feasible, we partition the system along the Navy's community structure. Each community consists of 500 to 5000 sailors of similar paygrade and job skills. One such community might be that of sonar technicians rated E7 through E9. A separate multi-agent job assignment system based on the IDA technology will be developed for each such community.

On the science side of the project we hope to implement a number of cognitive features of the IDA conceptual model that are not currently part of the running computational model. These include a transient episodic memory, the automization of action sequences, the involvement of emotions in decision-making and problem solving with novel problem. Mechanisms for all these have been designed and await resources for coding and testing.

Much additional information about the IDA project, including many technical details can be found at the link below in the form of online versions of published papers.

Link:
http://csrg.cs.memphis.edu/csrg/index.html

Please contact:
Stan Franklin
The University of Memphis, USA
E-mail: franklin@memphis.edu

IndiGolog: A Programming Language for Cognitive Agents

by Yves Lespérance and Hector J. Levesque

You have just arrived at the office. In a few minutes the office robot should bring you your morning coffee. It will have noticed your arrival from your logging in to the office's computer system. It will also bring you the airline tickets that have arrived in the mail. You got those with the help of your automated travel assistant agent…

In the not so distant future, we can imagine situations where people interact routinely with agents or robots that are fully automated and that make reasoned decisions for themselves about what to do. Likely this will first involve tasks that people find tedious such as filtering email, making travel arrangements on the web, or doing housework. To build an agent that can act intelligently and help a user carry out tasks like these, we need to provide the agent with reasoning capabilities that support building plans of action, explaining its observations and updating its knowledge, predicting the behavior of other agents, communicating and cooperating with them, etc.

For more than a decade, researchers from the Cognitive Robotics Group at the University of Toronto (under the leadership of the late Ray Reiter), together with collaborators at York University, and universities in Rome, Aachen, and elsewhere, have been developing a unified reasoning framework based on a simple logic of action, called the Situation Calculus, for specifying and implementing such intelligent agents. This approach involves first specifying in a declarative way the effects of the actions that can be performed in an application domain (eg, obtaining a coffee from a coffee machine) in this framework. Once this has been done, one can write complex high-level control programs for the agent in terms of these domain-dependent actions (eg, for delivering coffee to

Cognitive Robotics Group members.

people around the office). For this, a programming language called Golog was developed. Golog (which stands for Algol in Logic) looks a lot like a standard procedural programming language. But, the interpreter for this language uses the declarative specification to reason about what the effects of various possible actions would be. The programmer writes a program that can be sketchy and nondeterministic, and the interpreter automatically figures out a way to execute it successfully. The program can be quite vague, specifying only the goals to be achieved, or more detailed, specifying some steps and search control information, or fully detailed, specifying the agent's behavior completely. Later, an extended language called ConGolog was developed to support concurrent high-level programming. This facilitates the design of agents that have both reactive and proactive behaviors.

More recently, we have focused on extending the framework to support the design of agents that operate with incomplete information and acquire new information through sensing actions. A new extended programming language called IndiGolog was developed which allows the programmer to specify what parts of the program should be deliberated over (to find a plan). The program can perform some planning, then execute the plan and acquire new information, then do more planning, etc. The execution of a plan generated by a search block is automatically monitored and replanning is done when failures or changing conditions require it. We are currently extending the interpreter to support the generation of plans which themselves contain sensing actions. IndiGolog also supports a simple form of contingent planning where the dynamic environment is modeled as a simple deterministic reactive program. This is also being extended to deal with more complex nondeterministic environments. In parallel with this, extensions to the framework to deal with stochastic actions, uncertain knowledge, noisy sensors, temporal constraints, continuous processes, etc. are under development. Tools have also been developed to verify programs written in the framework.

Many robotics applications have been developed using our framework. At York University, a mail delivery robot application was implemented which optimized delivery routes and replanned routes when failures occurred or new orders arrived. At the University of Toronto a coffee delivery robot application that optimized temporal constraints was implemented. The framework was adapted for use in programming very inexpensive Lego Mindstorm® robots. Golog was also used as part of a very successful museum guide robot application developed at the University of Bonn in Germany, and later demonstrated at the Smithsonian in Washington. Our approach has influenced work at NASA on autonomous agents for spacecraft control. We have also been collaborating with MD Robotics (designers of the robot arms of the space shuttle and space station) on space vision applications.

Our framework has also been used to develop software agent applications, such as a banking assistant. An extension of Golog is currently being used at Stanford University for web service customization and composition.

Links:
http://www.cs.toronto.edu/cogrobo/
http://www.cs.yorku.ca/~lesperan/

Please contact:
Yves Lespérance
York University, Toronto, Canada
E-mail: lesperan@cs.yorku.ca

Hector J. Levesque
University of Toronto, Canada
E-mail: hector@cs.toronto.edu

From Isolated Components to Cognitive Systems

by Christian Balkenius

IKAROS is an infrastructure that will allow large-scale cognitive systems to be built by assimilating and combining results, methods and models from different research groups.

Before robots can move into our homes, offices and hospitals, a certain level of autonomy is needed. We need systems able to function in a human environment without causing too much annoyance - a high level of autonomy, but far less than you would need for, for instance, mission-critical operations in remote or hostile environments.

While it may be acceptable for a service robot to ask us whether we want fillet mignon or a hamburger for dinner, it will not be acceptable for a robot to ask for advice on how to open a door or how to go to the kitchen. In short, it will need the ability to solve all the small everyday problems that humans do not even have to think about. This type of everyday cognition relies as much on sensory processing and motor control as on high-level reasoning and planning. In fact, it uses all the subsystems studied within the field of cognitive science, ranging from emotion and motivation to attention and motor control. An understanding of the biological solutions in these areas may thus give important insights into the architecture that is required.

A common objection to biologically inspired systems is that "aeroplanes do not flap their wings", but it is important to realise that birds and jumbo-jets are solutions to two quite different design problems. When the design problems do coincide, such as in energy-efficient gliding flight, the design solutions are remarkably similar, such as the wing configuration of an albatross and a glider.

We want to suggest that the problem domain for the control systems of future autonomous robots is sufficiently similar to that of the human brain to merit a study of its biological counterpart. A domestic robot will need to exist and function in an environment that is remarkably, though unsurprisingly, well adapted to human cognition and motor behaviour. For humans, there is ideally no design problem; the environment is already adapted to us. For a robot, the best solution is to employ design strategies that are known to work in the environment, and those strategies are those that humans already use.

Today, a large number of computational models are available for many cognitive domains and brain regions involved in cognition. Although many of these models do not correctly reproduce all the functions of their corresponding brain areas, they are often sufficiently developed for use as components in larger systems. An important insight from our earlier research has been that many cognitive phenomena are system properties. The exact operations of the individual components are often not as important as how they are combined into a larger system.

The need for an infrastructure that would allow integration of models and methods from different cognitive research fields and different groups provided the motivation for the IKAROS project (named in tribute of a previous high-flying exercise in biological imitation). The project started in 2001 at Lund University Cognitive Science and builds upon our previous experience with simulators and control systems for robots. The goal is to develop tools that will allow full-scale cognitive architectures for future autonomous robots to be developed based on insights gained from human cognition.

The strength of IKAROS is its ability to assimilate different types of models and algorithms in such a way that they can be used as parts of larger systems. This makes the system open-ended, since support for new methods and hardware can be easily added. This is combined with support for communication between individual modules running on the same or different processors or on different computers on the Internet.

IKAROS systems are also inherently scalable. For example, the same architecture that runs on a small robot with a micro-controller and simple sensors can run on a large cluster computer with a full-scale vision system. Changing the inputs or outputs will automatically scale all individual components in the system to match. Another type of scalability is that architectures built with IKAROS can be moved seamlessly from simulation to real robots.

However, the most important aspect of IKAROS is that it is fully open and freely available. New versions are regularly published on the Internet and it runs on most computer platforms and operating systems. The first public version was released in February 2003. The current distribution already contains modules simulating many different brain regions, modules for image-processing and a collection of biologically motivated learning algorithms. The system also includes modules that communicate with external devices such as video cameras. An interface for mobile robots is currently under development.

We encourage other research groups to use IKAROS for their own research and to contribute new modules to the system. In the future, IKAROS will also include a database of benchmark problems to test and validate different cognitive architectures.

Link:
http://www.lucs.lu.se/IKAROS

Please contact:
Christian Balkenius and Jan Mørn
Lund University, Cognitive Science, Sweden
Tel: +46 46 222 32 51, +46 46 222 85 88
E-mail: christian.balkenius@lucs.lu.se, jan.morn@lucs.lu.se

Cognitive (Vision) Systems

by Henrik Christensen

'Cognitive Vision Systems' is a European project addressing issues related to categorisation, recognition, learning, interpretation and integration in relation to vision systems for intelligent embodied systems.

Over the last decade there has been significant progress in the fields of artificial intelligence, computer perception, machine learning and robotics. Yet there has not been major progress on truly cognitive systems. Cognition is here interpreted as 'generation of knowledge on the basis of perception, reasoning, learning and prior models'. Cognition is not a *passive process*, where an observer merely is monitoring an external environment. In addition the system has facilities for communication with the environment through which it also can articulate its knowledge. The system is embedded in the world and interacts with its environment to gather knowledge and perform its mission. Consequently, a *cognitive system needs to be embodied* and it has a number of tasks that define its mission objectives. The operation of the system is defined by its mission objectives and it acquires knowledge about the environment through its perception system. The perception is active in the sense that the system can use the embodiment to interact with the environment, which can be used to change the state of the environment, which subsequently can be observed by the perception system. The sheer amount of information available in the external environment calls for methods to generate 'abstract' models.

One important part of generation of models is the ability not only to RE-cognize objects, but also to perform recognition by means of categorisation. Categories form an important component to allow management of an abundance of information, association of function, physical layout, etc. for objects, situations and events. Categorical perception is, however, a major challenge as is seen on in the adjacent image. A task could here be "count the number of chairs". This is a non-trivial task. Without understanding of the physical layout of the scene some objects may be recognized incorrectly and using contours only, some of the shadows might be confusing. In addition the pictures on the wall do not represent places to sit. Visual cognition is obviously only one of many potential modalities that are of interest to cognitive systems. In addition for some objects there might be a need to interact with the environment to determine if an object qualifies for a particular task. For instance, is a chair stable enough to allow sitting down? Some of these qualities can only be determined by interaction.

Another fundamental component of any cognitive system is memory. Inherently memory is limited and there is a need to consider how memory can be utilized for different purposes: context information, spatial layout, abstraction, etc. Memory is, however, limited which call for efficient methods to manage this scarce resource. This requires use of attention mechanisms for select information of interest, and data-mining methods or machine learning to generate abstractions and derived representations. A basic quality of memory is also forgetting or intelligent garbage collection methods.

The relation between prior models or long(er)-term memory and working memory is another issue to be addressed. As part of the study of memory there is also a need to study pedagogic models for acquisition of information. What is

Example image that illustrates the complexity of recognizing chairs as no single technique in terms of contours, appearance, components etc. is adequate to correctly allow 'counting of the number of chairs'. (Source: Bülthoff, Max Planck Institute for Biological Cybernetics (MPIK), Tübingen, Germany.)

the best possible model to interact with a system to allow it to acquire a new skill, a task model or the representation of a new concept? Can traditional pedagogic models also be used to teach artefacts new concepts in an efficient manner? Cognitive science has proposed models for how infants acquire new models of the world; one such example is the learning models proposed by Piaget. Can such models be directly applied (can we endow the system with curiosity) or can be rephrase these methods so as to make them operational and applicable for artificial cognitive systems.

The concepts outlined above are fundamental questions addressed in the EU project 'cognitive vision systems' in which issues related to categorisation, recognition, learning, interpretation and integration are addressed in relation to vision systems for intelligent embodied systems. For categorisation of objects a new hybrid model that integrates multi-models for recognition with different types of memory has been proposed. In addition the relation between spatial models (of objects) is being integrated with models for scene dynamics to capture episodic information and tie these to particular objects. The relation between reasoning, interpretation, recognition and processing of basic information cues is another fundamental problem studied. How does context allow control of the visual process to make it tractable, while at the same time allowing for enough richness to detect unexpected events? Traditional formal models in AI have not had enough richness, but recent progress on reasoning under uncertainty shows promise in terms of richness and efficiency. Finally a Piaget inspired model of skill and task acquisition is being implemented as a basis for teaching robot like 'creatures' to interact with the environment. At present basic technologies are available and pair wise integration of techniques is performed to allow in-depth studies of the interaction between vision, AI, cognition and biology, psychology, computer science, and robotics.

Link:
http://cogvis.nada.kth.se

Please contact:
Henrik I. Christensen,
KTH Royal Institute of Technology, Sweden
E-mail: hic@nada.kth.se

Retrieval of Images based on Visual Content: A Biologically Inspired Multi-Agent Architecture

by Socrates Dimitriadis, Kostas Marias and Stelios Orphanoudakis

Understanding the mechanisms of visual perception is important in a number of application areas, including content-based image retrieval (CBIR). An interdisciplinary approach to CBIR, in the framework of cognitive science, is expected to yield significant results where traditional approaches have failed to do so. In the Computational Vision and Robotics Laboratory (CVRL) of the Institute of Computer Science - FORTH, an ongoing research effort aims to develop and implement an experimental platform for the investigation of CBIR, based on a biologically inspired multi-agent architecture.

In recent years, interdisciplinary research has produced significant results in a number of areas in which traditional approaches appeared to have reached their limit. Based on the joint effort of biological neuroscientists, computational neuroscientists, psychologists, and computer scientists, the interdisciplinary field of cognitive science has emerged as the field of research in which brain function is investigated with respect to behaviour and mental activity. Visual perception is a major brain activity and has attracted the attention of many cognitive scientists. Understanding the mechanisms of visual perception is important in a number of application areas, including content-based image retrieval (CBIR).

In the Computational Vision and Robotics Laboratory (CVRL) of ICS-FORTH, an ongoing research effort is aims to develop and implement a biologically inspired image retrieval methodology. Due to the fact that images represent a particularly large volume of information, the efficient and possibly intelligent browsing of images based on visual content is becoming increasingly important in application fields which make use of large image databases, such as diagnostic medical imaging, remote sensing, entertainment, etc. Surfing on the Web may also be facilitated and made more intelligent if visual browsing is exploited.

Humans undoubtedly possess the ability to process visual information efficiently and to identify images as being similar based on their visual content. However, computational approaches currently fall short of matching this ability. The reasons for this are many and cannot be explored in this short article. It is our research objective at ICS-FORTH to develop and implement CBIR mechanisms that are perceptually relevant and based on a biologically inspired system architecture. We must therefore take into account what is currently known about how humans process visual stimuli and subsequently represent and store visual information, so that it may later be retrieved efficiently from visual memory.

It is known that the human cerebral cortex consists of many cortical areas, each one specialised and able to execute a specific task. The visual cortex consists

of a set of specialised modules that are responsible for the processing of a specific visual feature and can execute their respective tasks independently and in parallel. An image similarity decision is not based on any one of these processes acting in isolation, but is rather the result of the integration of their individual contributions. This biological paradigm resembles the computational approach adopted in multi-agent systems. Exploiting parallelism in complex dynamic environments with a high level of uncertainty is a key ingredient of multi-agent computational systems that model biological processes and behaviours.

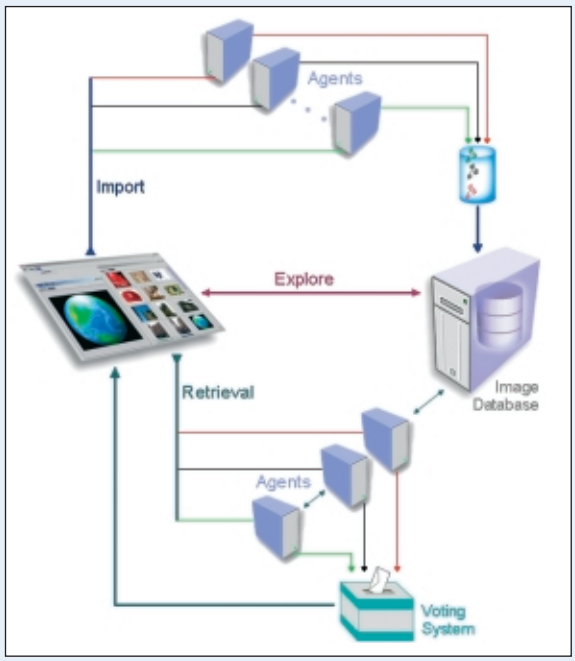

Architecture of the content-based image retrieval system.

A biologically inspired multi-agent architecture has been developed at ICS-FORTH, and provides a platform both for the experimental investigation of a variety of problems arising in CBIR and for the assessment of the perceptual relevance of similarity retrieval schemes involving the cooperation of selected agents. As shown in the figure below, this architecture consists of four main components: a set of software agents, an image database, a graphical user interface, and a voting system that supports a set of alternative voting schemes. Images can be imported into the database, viewed individually or in galleries, and retrieved based on their visual content. Images that are similar to a query image may be retrieved from a specific set of candidate images (gallery) or from all images contained in the database. The query image and the set of candidate images are both selected by the user. Each agent works independently and computes the similarity between the query image and each candidate image, based on its specialised similarity criterion and matching algorithm. Integration of partial results obtained by different agents is achieved through the voting system, using a voting scheme selected by the user. The system loads and handles dynamically the set of available agents and voting methods. Furthermore, the reasoning on which the final ranking of candidate image similarity is based is exported, thus allowing the user to interpret the response to a specific query and to fine-tune relevant parameters in order to improve the response to subsequent queries. This reasoning consists of the contribution of each agent to the final ranking of each candidate image, according to the voting scheme selected by the user.

The CBIR system is fully scalable and can easily be extended with additional agents or voting schemes in order to take into account the specific requirements of different experiments and the class of images under consideration. It has been developed in JavaTM and makes use of the Java Advanced Imaging package.

Please contact:
Stelios Orphanoudakis, ICS-FORTH
Tel: +30 2810 391605
E-mail: orphanou@ics.forth.gr

The Cogitoid: Towards a Computational Model of the Mind

by Jirí Wiedermann and Martin Beran

The cogitoid is an algorithmic model of the cognitive processes occurring in the mind of living organisms. A cogitoid presents a multimodal interactive learning algorithm that learns from experience. Thanks to its formal definition, the cogitoid allows formal reasoning about its ability to perform fundamental cognitive tasks. At the conceptual level, the cogitoid also offers a plausible framework for explaining higher mental functions.

In the mid-nineties, a project in computational mind modelling was set up in the Institute of Computer Science at the Academy of Sciences of the Czech Republic. The aim of this project has been to develop a basis for an algorithmic theory of the mind by designing a formal model of the mind, deriving results about its cognitive abilities, investigating its explanatory potential regarding higher mental processes and, last but not least, experimentally verifying its viability in solving concrete cognitive tasks.

Computationally, a cogitoid is viewed as a finite transducer that transforms a potentially infinite sequence of its

sensory inputs into a similar sequence of actions. A cogitoid is realised as a network of concepts with associations among them. In the interaction of a cogitoid with the environment, both new concepts and associations are formed, strengthened or weakened. Emotions are also built into the model with the help of special predefined concepts, the activation of which modulates the sequence of activation of other concepts. A cogitoid works by iterating the cycle sense-compute-act.

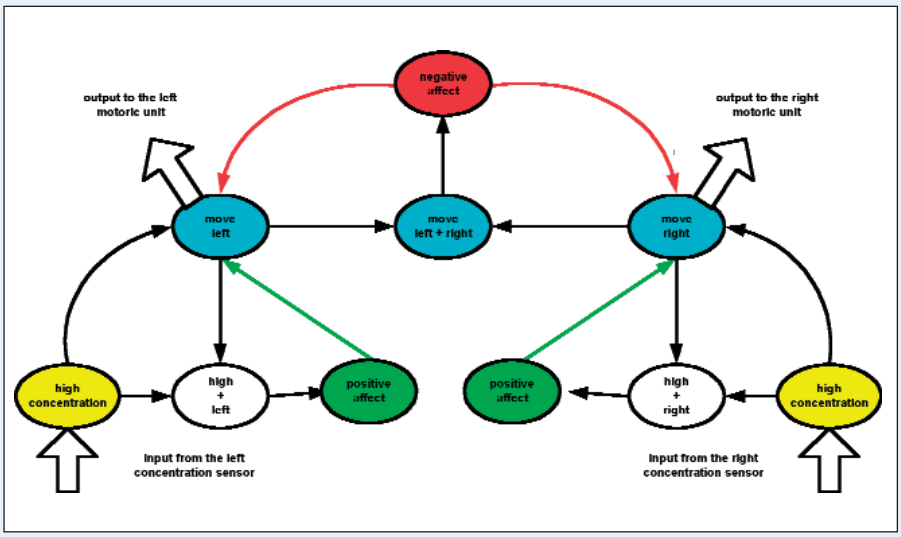

A scheme of a cogitoid.

For such a model, theorems showing the ability to learn a sequence of stimuli (simple conditioning), learning by similarity, Pavlovian conditioning, and operant conditioning (learning by punishment and reward) have been proven. Proving theorems about more complex behaviour is complicated by the fact that in a cogitoid, many of the interactions that emerge among concepts and associations are difficult to capture formally. Therefore, the higher cognitive functions are not described by theorems, but rather by claims with supporting reasoning and verifying experiments. For cogitoids with sufficiently many concepts and rich interaction with a feedback from the environment, one can expect automatic formation of concept clusters in which concepts related to certain activities, contexts, or subjects are assembled. The clusters represent a basis for the mechanism of understanding and intention formation. By navigating among the respective clusters along the frequently strengthened associations, a behaviour based on the experience unfolds itself. In this way a number of high-level mental phenomena, such as imitative learning, self-control, the formation of the concept of self, language acquisition and generation, thinking, and a rudimentary form of consciousness, can be explained.

Recently there were several computational experiments performed with simple cogitoids. In the simplest case, a so-called 'computational bacterium' driven by a cogitoid has been designed, which was able to learn to move in one dimension towards a higher concentration of nutrients. A scheme of the respective cogitoid that eventually developed after interaction with a changing environment is depicted in the accompanying picture. This experiment has shown a need for two additional conditions to be defined for a cogitoid to operate successfully. The first condition deals with the question of a proper setting of positive and negative effects ('emotions') in order to strengthen desired and suppress unwanted behaviour. The second condition is concerned with the problem of the initial activity of a cogitoid that has not yet learned any behavioural patterns. In such cases, it appeared useful to equip the cogitoid with a body, which is driven by a very simple heuristic algorithm simulating 'reflexes'. The cogitoid then 'observes' the activities of its own body and after some time it takes over the control of its body. This has been verified in a more complex experiment where an originally blind agent was able to learn to use its additionally provided vision ability for avoiding collisions. A Masters thesis aiming at programming a cogitoid on a cluster of PCs has also been initiated. All these experiments have been performed at the Department of Computer Science in the Faculty of Mathematics and Physics at Charles University, Prague.

Currently, experiments that are more complex are still under development. These will help in further understanding the abilities and limitations of cogitoids in realising non-trivial cognitive tasks. They will also help in finding better ways of designing a cogitoid as well as 'flaws' in its definition. The genetic algorithms for automatic cogitoid evolution and tuning are also under consideration. Finally, yet no less crucially, the plausible explanation of further details of the algorithmic mechanism behind the higher brain functions remains a focus of the current research. The research reported has been partially supported by GA CR grant No. 201/02/1456.

Links:
See the related technical reports at http://www.cs.cas.cz/ and http://www.ms.mff.cuni.cz/~beran

Please contact:
Jirí Wiedermann, Institute of Computer Science, Academy of Sciences of the Czech Republic/CRCIM
Tel: +4202 6605 3520
E-mail: jiri.wiedermann@cs.cas.cz

Modelling the 'Homunculus'

by András Lörincz

The Neural Information Processing Group of the Eötvös Loránd University Budapest has been engaged in researching reflexive systems, capable of collecting experiences, including reinforcement. The system can learn from experiences and can accept directions. The system envisioned will keep an eye on itself and may seem to know what it is doing.

Generally speaking, the processing of signals that convey information can be considered as a transformation into another form that still carries the whole or just a piece of the original information. The environment feeds the system with some inputs and the system output represents (a part of) the environment. Whilst most models address the problem of coding inputs and making efficient internal representations, we are more concerned about the fundamental problem of making sense of these representations. In our view, the central issue of making sense or meaning is to provide answers to questions like "what does it mean?" in terms of our past experiences, or "how is it related?" in terms of known facts. In other words, making sense is inherently related to declarative memory. As a consequence, the homunculus (the little strange person who sits in the 'Cartesian theater' - to use the words of Dennett) becomes a central issue. There is a related fallacy that says that the internal representation is meaningless without an interpreter. This fallacy claims that all levels of abstraction require at least one further level to become the corresponding interpreter. Unfortunately, the interpretation - according to the fallacy - is just a new transformation and we are trapped in an endless regression. This problem could be more than a philosophical issue. We are afraid that any model of declarative memory or a model of structures playing a role in the formation of declarative memory could be questioned by the kind of arguments provided by the fallacy.

It is known that human-computer interface needs 'mind-reading' capabilities, eg, (i) the recognition of different emotional states, (ii) the actual state of thinking, such as concentrating, being tired or lost, etc. We are currently incorporating our reflexive architecture into human-computer interfaces, which can collect information about the 'state' of the user by measuring head and eye movements as well as facial expressions. The system will be capable of responding by animating facial expressions. The computer will use adaptive probabilistic reflexive architectures. The computer will make use of a visual thesaurus. For example, the figure depicts one frame of a short clip expressing 'wondering'. The interacting computer will use, interpolate, morph, and combine a collection of such short clips. Any of these clips is an 'action', from the point of view of an interacting computer, whereas the response of the user can form the reinforcing feedback. This is the basic scheme of reinforcement based optimization of human-computer interaction, serving the user. Note that this is the basic scheme that we use in our everyday interaction from early childhood when we 'live together'.

In our theoretical efforts, we start by claiming that the paradox stems from vaguely described procedures of 'making sense'. The fallacy arises when we say that the internal representation should make sense. To the best of our knowledge, this formulation of the fallacy has not been questioned except in our previous work, where the fallacy was turned upside down by changing the roles: Not the internal representation but the input should make sense. The proposal is the following. The input makes sense if the same (or similar) inputs have been experienced before and if the input can be derived or regenerated by means of the internal representation. According to this approach the internal representation interprets the input by (re)constructing it.

The idea behind this approach is to keep the infinite regression, make it a converging regression and execute this regression in a finite architecture. The change of the roles gives rise to a reconstructing loop structure. The loop has two constituents; the top and the bottom. The top contains the internal representation. The representation generates the reconstructed input. The bottom computes the difference between the actual input and the reconstructed input. This difference, the reconstructed error, corrects the internal representation, and so on. This is a finite architecture with a converging, but - in principle - endless iteration. In turn, the infinite regression is transformed into a converging iteration, and a route has been opened to reduce the fallacy to the problems of stability and convergence.

The approach - in principle - leads to a reflexive system, which can infer about

the external world, about its own state and about the actual interaction between those. Notably, there are relatively strong (mathematical and computational) constraints on how such a reconstruction network should work. Constraints emerge, for example, from optimization of information transfer between bottom and top. Such constraints severely restrict modelling.

We accepted these constraints and started to develop a model of the 'homunculus'. The uniqueness of our model is that starting from a relatively small set of hypotheses, many structural and functional features can be derived. These features are indirect predictions of the model. Without including biological constraints beforehand, we could show the emergence of some specific low order memory functions. The emerging properties of our model:

- explains general properties of the brain, such as priming and repetition suppression
- recognizes novelty before searching the database
- has a correct order of learning
- explains properties of certain brain diseases
- provides a unified view of control and sensory processing
- is built of elements, which have a local Kalman-filter like structure with proper Hebbian-learning rules
- encompasses reinforcement learning in a natural fashion
- finds domains, which are almost deterministic
- boosts prediction and goal-oriented planning
- has striking similarities with known brain structures
- provides a view for consciousness and introspection
- has a philosophical ground, which goes back to the works of Locke and Hegel.

To give an example, consider novelty detection. Several works have shown that maximization of information transfer gives rise to sparse representation for natural stimuli. Then novel information (which could be a novel natural stimulus) gives rise to a non-sparse representation. In turn, the distribution of neural activities at the top indicates whether the actual input is familiar or novel: distinctions between novelty and familiarity can be immediate upon (the first) bottom-up information transfer in the loop.

Link:
http://people.inf.elte.hu/lorincz/pub.html

Please contact:
András Lörincz, Eötvös Loránd University, Budapest, Hungary
Tel: +36 1 209 0555 / 8473
E-mail: lorincz@inf.elte.hu

Emergence of Representations through Interactions of a Robot with the Real World

by Nicolas P. Rougier and Frédéric Alexandre

Research by the Cortex Research Team at INRIA Lorraine is focusing on the design of biologically inspired models with numerical and robust mechanisms, which can promote the emergence of multimodal representations within the framework of real-world robotics.

We have been studying the various kinds of information representation that may be useful to an autonomous robot seen as a cognitive system. These studies integrate data from the neurosciences into computational models, and in this framework, we have demonstrated how procedural, declarative and working memories can emerge from specific neuronal structures, and can become critical at different levels of information processing.

For instance, in a simple obstacle avoidance system, a neural network model may only learn the required sensor-motor coordination between sensors and actuators. This is called procedural memory, and is one of the roles of the associative cortex. In a more complex task requiring an autonomous system to identify certain locations in an environment (eg because of obstacles), the model needs to memorise descriptions (even very basic ones) to be able to reach these locations. This is called declarative memory, and is one of the roles of the hippocampus. In an even more complex task, a robot may need to perform a sequence of tasks in which order is critical, and therefore needs a system of representation that allows it to recall which tasks have already been performed, and which remain to be done. This is called working memory, and is one of the roles of the prefrontal cortex.

We have been exploring these three kinds of representations using computational models and have demonstrated the basic principles for each of them. Nonetheless, our studies of these systems at the global level (the level of the whole model) have, for several reasons, also underlined the need to re-think the computational mechanisms implied at the unit level. On the one hand, neuronal models with such specific memorisation abilities need to collaborate and exchange information during real robotic tasks requiring the coordination of these properties. On the other hand, using these models with real robots in the real world requires the design of very robust numerical models.

As a consequence, we are currently working on the design of robust numerical mechanisms that can promote the emergence of these different kinds of information representation. More precisely, we have been studying a model of the cerebral cortex (Burnod, 1989) that emphasises the role of units,

which can be described as assemblies of neurons, allowing them to process information at a higher level than the classical neural unit. We have demonstrated in the past that networks of such units are greatly advantageous for information processing, since they allow the generalisation of knowledge from a small set of experiences. Furthermore, we are now focusing on numerical mechanisms à la CNFT. These deal with a restrained connectivity between units, which is quite critical in enabling maps to interact with each other (eg visuo-motor maps, visuo-auditory maps, polymodal maps).

In the framework of a French CNRS Robea project, we are first developing biologically inspired and robust visual and motor representations together with multimodal integration, which can endow a robot with the ability to predict the consequences of its actions on the real world (an arm with nine degrees of freedom).

In the framework of the FET European project MirrorBot, we are interested in more abstract representational learning. We have defined some complex behavioural protocols for our robot (object localisation, recognition and reaching, action planning and imitation, speech recognition and production) and we are observing how these memorisation mechanisms interact with each other to allow multimodal learning. We are particularly interested in a better understanding of the role of mirror neurons, originally observed by G. Rizzolati and his laboratory, a member of this project. These neurons have been found to be activated in a monkey's prefrontal cortex by the observation of an action performed by the monkey itself, by another monkey or even by a human. This cortical region corresponds to a region devoted to language in the human brain, and is thought to be the basis for complex abstract representations of action.

Through this ongoing research into multimodal integration and autonomous robot navigation, we aim to show that useful representations for world understanding and language acquisition can emerge from interactions with the world, and that building such representations can be performed by robust numerical mechanisms inspired by neuroscience.

© INRIA / Photo A.Eidelman

The Koala robot and its pan/tilt camera is used to test the robustness of the models in real world conditions.

Link:
http://www.inria.fr/recherche/equipes/cortex.en.html

Please contact :
Nicolas P. Rougier, LORIA-INRIA
Tel: +33 3 83 59 30 92
E-mail: Nicolas.Rougier@loria.fr

Causal Structures in Embodied Systems

by George Kampis and László Gulyás

Embodied cognition is the front line of cognitive research. The Cognitive Science Program at Eötvös Loránd University in Budapest is running a project in cooperation with SZTAKI and other institutions. In its current phase, the project aims to develop simulations based on causal principles that unify autonomous robots, evolutionary processes and cognitive systems, thereby testing basic principles of open-ended real-world intelligence through interaction.

In the last decade it has become generally clear that human and animal cognition is based on embodied functioning rather than abstract representation. To operate an embodied system, sensory-motor loops, learning, and situated actions coordinated by low-level control are necessary, rather than reasoning, planning, or knowledge. This recognition has led to a shift in emphasis in cognitive systems research. Instead of representation, situatedness and environmental action are believed to supply essential information for the agent. The new approach is supported by findings in developmental child psychology and cognitive linguistics, and by the dynamic view of cognition.

Most of the current models developed in this new spirit are anti-representational in the same sense as old cybernetics, where every material function was reduced to feedback mechanisms and dynamic control. In the embodied framework, research on genuinely cognitive structures is neglected. An additional problem is that embodiment tends to invite a controversial focus on internal experience, a factor that is not accessible to scientific modelling.

Our research started from the hypothesis that there exist intermediate cognitive structures between situated behaviour and mental experience, and that through these we can get a handle on cognitive functions. Recent interest in systems where motor control, neural learning, vision and symbol systems meet can supply a similar assumption. The paradigm of embodied cognition must

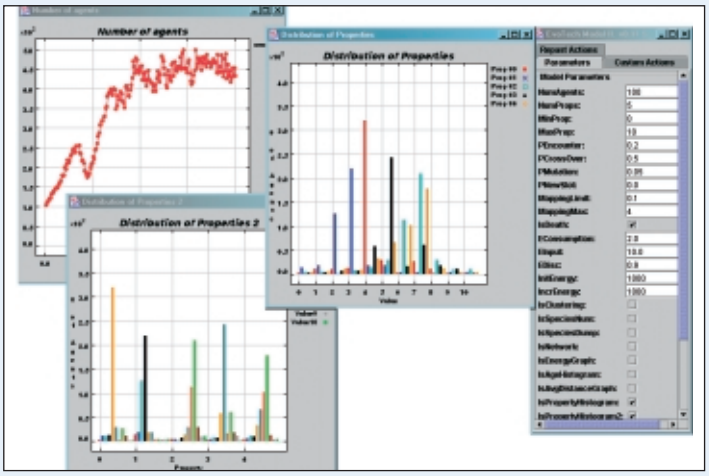

Figure 1. Part of the frameset of the simulation model. It shows population numbers, surviving phenotypes, and their distribution.

switch its focus of inquiry from lower-level, insect-like intelligence, which has dominated earlier research, to vertebrate cognition. The area in which higher organisms are superior to lower organisms is the integrated use of their various faculties. We are looking at how this may happen.

The Primacy of Causality

The central question for an embodied cognitive faculty is that of causal interaction, both at the level of bodily events and at that of mental models. When looking for the meaning of causality for a real-world body, the usual conception of causation needs to be extended. Causation is usually assumed to be a relationship between events. This is a simplification, which has its origin in the abstract representational conception. A realised action involves simultaneous activities that take place in several parallel layers. In other words, causality has 'depth'. One form in which causal depth can be studied is via supervenient levels acting permanently together. Our interest lies in more dynamic ways in which the various related causal modes can influence each other.

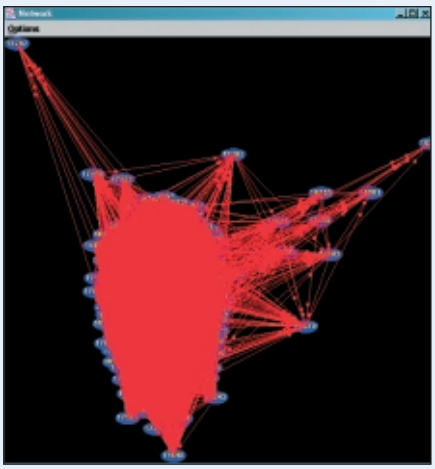

Figure 2: Evolved, stable species with escaping mutants; red lines denote mutual reproduction capability.

An example is the orchestrated use of different sensory modalities in active perception where, for instance, a visually guided mental model can recursively enact actions that approach the signal and lead to new percepts that help in stabilising the mental model. Here the underlying concept is that of a complex body with context-dependent causal faculties, as opposed to a simple one that receives externally fixed inputs.

An Application: Evolutionary Technology

The example chosen to illuminate the framework is evolution. To free ourselves from biological details we consider evolutionary technology instead of natural evolution. The ultimate goal is to develop artificial organisms that perform increasingly complicated tasks. We are currently in the process of developing a simple test-bed for the causal principles.

Evolution and cognition have striking parallels. In both cases it is complex bodily properties of a physically realised agent that collectively determine a historical process of structure formation, which we often view from the perspective of its end results. If we want to understand the process of origin of the inner structure and the role that real-world causality plays in it, evolution is a good starting point. Evolution is also a useful example for applying causal concepts of embodiment because equipping abstract systems with causal powers may help us to understand what makes an evolutionary process work. Selection is only half of the answer, as was recently demonstrated by studies on robotic as well as software-based systems, which failed to improve beyond a certain point.

The same causal principle that underlies active perception can also assist selection. The assumption is that an adaptation that produces a new phenotype can also switch interactions between organisms to a different causal mode, which makes a new adaptive process possible.

We developed an agent-based simulation model using the RePast package. Organisms are agents that selectively feed, reproduce and die, based on their phenotypic properties described in variable length records. As adaptation progresses, new property sets extend the records, and as a result, selection can spontaneously switch between the defining properties of an interaction. The aim is to develop functionally disjoint subpopulations specialised for the use of different property sets. The first results have recently been reported, showing the possibility of progressive evolution productive of new selection effects, as an illustration for the causal principles of embodiment.

Some steps of the ongoing research outlined in this article were carried out while George Kampis was Fujitsu Visiting Associate Professor at the Japan Advanced Institute of Science and Technology and during the visit of László Gulyás at Harvard's Center for Basic Research in the Social Sciences (CBRSS).

Links

http://hps.elte.hu/~kampis/projects/EvoTech.html

http://www.jaist.ac.jp/~g-kampis/

http://www.sztaki.hu/~gulyas/indexE.html

http://www.cbrss.harvard.edu/people/gulyas.htm

http://www.kzoo.edu/physics/ccss/research.html

http://repast.sourceforge.net/

Please contact:
George Kampis, Eötvös Loránd University, Budapest, Hungary
Tel: +36 1 372 2925
E-mail: gk@hps.elte.hu

Swarm-Bots: Swarm of Mobile Robots able to Self-Assemble and Self-Organize

by Stefano Nolfi, Jean-Louis Denebourg, Dario Floreano, Luca Gambardella, Francesco Mondada and Marco Dorigo

Swarm-bots are a collection of mobile robots able to self-assemble and to self-organize in order to solve problems that cannot be solved by a single robot. These robots combine the power of swarm intelligence with the flexibility of self-reconfiguration as aggregate swarm-bots can dynamically change their structure to match environmental variations.

SWARM-BOTS, a project funded by the Future and Emerging Technologies program of the European Community (project IST-2000-31010), focuses on the design and the implementation of self-organising and self-assembling biologically-inspired robots. The project benefits from the joint activity of four European research institutes: the Artificial Intelligence Laboratory of the Université Libre de Bruxelles (IRIDIA); Dalle Molle Institute of Artificial Intelligence Studies of Lugano (IDSIA); the Autonomous Systems Laboratory of the Swiss Federal Institute of Technology (ASL-EPFL), and the Institute of Cognitive Sciences and Technology of the Italian National Research Council (ISTC-CNR).

A swarm-bot is an aggregate of s-bots (mobile robots able to self-assemble by connecting/disconnecting from each other) that can explore, navigate and transport heavy objects on rough terrains in situations in which a single s-bot would have major problems to achieve the task alone (see Figures 1 and 2). Each s-bot has simple sensors and motors, limited computational capabilities, and physical links that allow it to connect to other s-bots. The swarm-bot should move as a whole and reconfigure along the way when needed. For example, it might adopt a different shape in order to go through a narrow passage or overcome an obstacle. An aggregate should form as a result of self-organizing rules followed by each individual s-bot rather than via a global template. Thus swarm-bots combine the power of swarm intelligence, as they are based on the emergent collective intelligence of groups of robots, and the flexibility of self-reconfiguration as aggregate swarm-bots can dynamically change their structure to match environmental variations.

We have now completed the hardware design (the first hardware prototypes will be ready within the next few weeks). As shown in Figure 1, each s-bot has a cylindrical body with a diameter of 116mm and consists of a mobile base provided with two differential drive mechanisms controlling tracks and wheels, and a main body with two grippers that allow it to assemble with other s-bots or to grasp objects. The first gripper is supported by a mobile structure that can rotate around a horizontal axis and the second gripper is supported by a motorized arm that allows large movements through the vertical and horizontal axes. The main body rotates with respect to the base supporting the tracks. From the sensory point of view, each s-bot has proximity sensors, light sensors, accelerometers, humidity sensors, sound sensors, an omni-directional color camera, light barrier sensors (on the grippers), force sensors etc.

We are following two different but complementary research directions in order to develop the control systems of a swarm-bot. One consists in building control systems that mimic the characteristics of biological systems such as social insects. The other consists in building control systems that are only loosely inspired by what we know about real organisms and develop their ability through a self-organization process based on artificial evolution. By following the first approach, we have been able to develop, for example, the

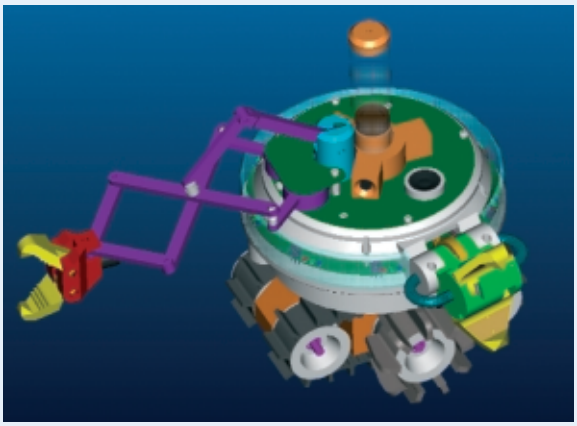

Figure 1:. The design of a single s-bot.

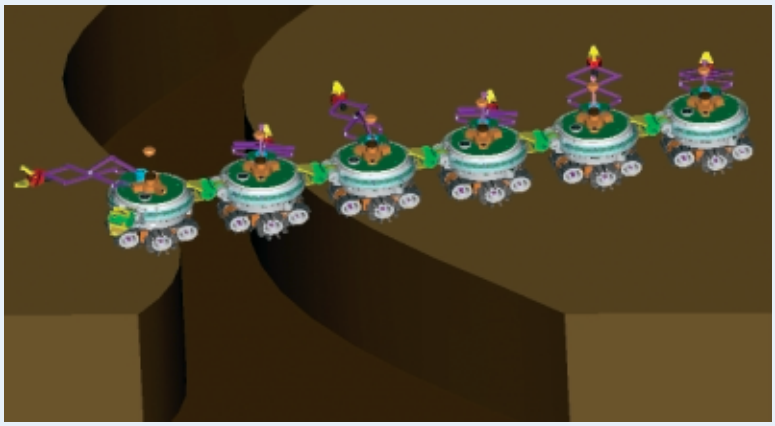

Figure 2: A swarm-bot made of several assembled s-bots passing a fosse.

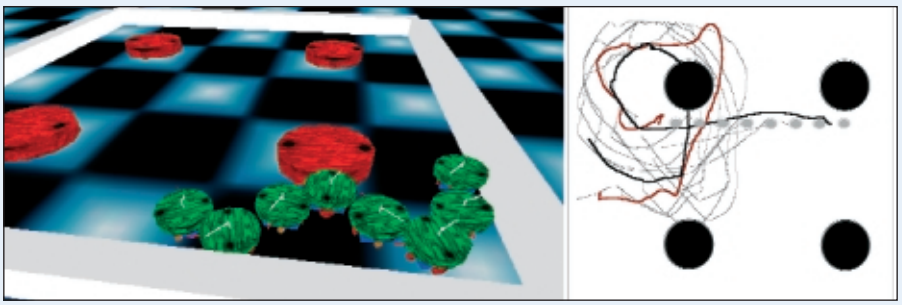

Figure 3: A swarm-bot displaying collective obstacle avoidance.

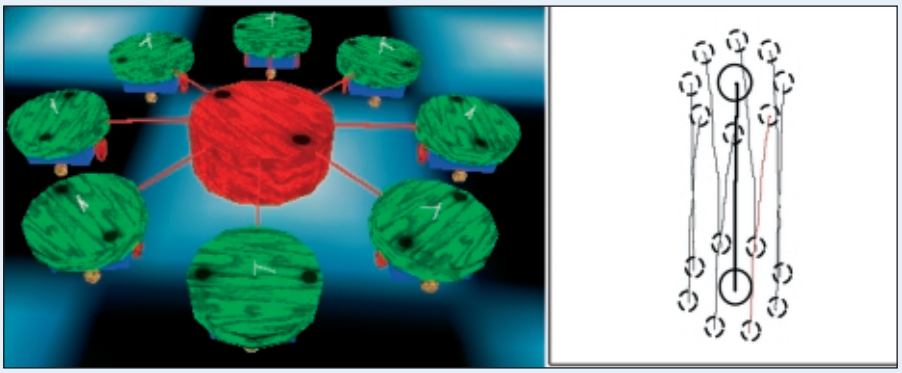

Figure 4: A swarm-bot displaying object pushing/pulling.

control systems of a group of simple LEGO™ robots that are able to aggregate into chain formations. By following the latter, we were able to develop simulated s-bots displaying coordinated movements, collective obstacle avoidance, and object pushing/pulling (see Figures 3 and 4).

Figure 3 shows a swarm-bot displaying collective obstacle avoidance. On the left, we have eight simulated s-bots assembled into a snake-like formation placed in an arena surrounded by walls with four cylindrical obstacles. On the right, we see the trajectory produced by the swarm-bots that move in the arena. The small gray circles represent the initial position and shape of the swarm-bots. The square and the large full circles represent the walls and the obstacles. The lines show the trajectory of the s-bots. As can be seen, the swarm-bot is able to avoid obstacles and pass narrow passages eventually deforming its shape according to the configuration of the obstacles.

Figure 4 shows a swarm-bot displaying object pushing/pulling. The swarm-bot consists of 8 s-bots connected to an object. On the right can be seen traces made by the s-bots (thin lines) and the object (bold line) during a few seconds. Large and small dotted circles represent the initial (bottom) and final (top) positions of the object and of the s-bot, respectively.

Potential applications of this novel type of robot are, for instance, semi-automatic space exploration, rescue searches or underwater exploration.

Links:
http://www.swarm-bots.org/

Please contact:
Marco Dorigo
Université Libre de Bruxelles
Tel: +32 (0)2 6503169
E-mail: mdorigo@ulb.ac.be

Cognitive Soccer Robots

by Alessandro Farinelli, Giorgio Grisetti, Luca Iocchi and Daniele Nardi

The effective development of a system consisting of autonomous heterogeneous robots that have to perform complex tasks over long periods of time, operating in a partially known and partially observable environment, implies addressing a wide range of issues. These issues involve many areas of Artificial Intelligence, from computer vision to deductive systems.

Since 1998 we have been developing autonomous robotic systems that have been used within the RoboCup competitions. RoboCup is an organization that promotes research in the field of autonomous robotics by organizing annual world competitions of soccer-player robots. Such competitions are made within different leagues depending on the type and size of the robots. We have created two robot teams: SPQR-Wheeled and SPQR-Legged belonging to different leagues, and although the robots composing the teams are quite different (various kinds of wheeled robots sized about 50x50 cm for the Wheeled team, and Sony AIBO four legged robots for the Legged one), they share the same cognitive system.

The main features required for the design and the implementation of a team of cognitive robots are: (i) a framework for defining the robots' knowledge of the environment and its capabilities; (ii) a reasoning system able to derive plans (ie programs) that must be executed to achieve a given goal; (iii) a plan execution mechanism; (iv) a coordination module that is in charge of assigning sub tasks to each robot in the team in order to achieve a global team goal.

Architecture
The cognitive robotic system, which we have implemented on several robotic platforms, is based on a hybrid layered architecture, with two levels: the Operative Level, in which the information is handled numerically, and the Deliberative Level, which relies on a symbolic representation of the robot's knowledge. In particular, at the Deliberative Level, the environment is described by a knowledge base containing both axioms and static facts (i.e background knowledge) provided by the designer, as well as dynamic facts (ie symbolic information) corresponding to the data acquired through the robot's sensor during its mission.

Domain and Plan Representation
Our planning and reasoning framework relies on an epistemic representation of the world, which explicitly takes into account the robot's knowledge, rather than the true world status. The domain description includes the specification of the basic actions the robot can perform, which are defined in term of preconditions (conditions that must be verified for an action to be executed) and effects (conditions that are verified after the execution of the action). Our framework makes it possible to define two kinds of actions: ordinary actions - which effect changes in the world, and sensing actions - which effect changes in the robot's knowledge (ie in its mental state) and allows for sensing the value of a property that is not known.

Plans are represented as transition graphs, whose nodes represent epistemic states and whose edges represent actions causing state transitions. The graph representing a plan has an initial node from which the execution starts and a set of final states in which the goal is satisfied. The execution of an ordinary action causes a deterministic state transition, while the execution of a sensing action causes a state transition depending on the sensing outcome. Moreover, since in the real world the execution of an action is not guaranteed to succeed, a set of recovery states are associated with each action.

Plan Generation and Execution
We have defined a logic formalism and developed a plan generator that, given: i) the domain specification as a set of axioms; ii) the actions specification, in terms of pre-conditions and effects; iii) an initial state; iv) a goal, generates a plan for reaching the goal from the initial state. Through the use of sensing actions, the plan generated can contain both if and goto constructs. It is also possible to manually edit plans using a graphic tool.

A plan is executed by navigating the plan graph, starting from the initial node and terminating when a goal node is reached or when the plan fails. A plan failure can be caused by the failure of some of the execution conditions of the action, and will activate a plan recovery or a new plan selection depending on the presence of edges exiting from the recovery state.

Coordination
Coordination among robots is achieved by a coordination module that selects which sub-task must be accomplished by every robot. In our cognitive systems, each task is related to a goal to be achieved by a robot and thus to a plan to be executed.

Tasks are assigned during robot operation by means of a distributed protocol that allows information about the current state of each robot in the team to be shared, and provides the necessary autonomy in case of network failures. The state of each robot and its capabilities are evaluated in order to decide to which robot a given sub-task should be assigned and a broadcast agreement among all the robots is performed in order to effectively assign sub-tasks to robots, while avoiding interferences among them.

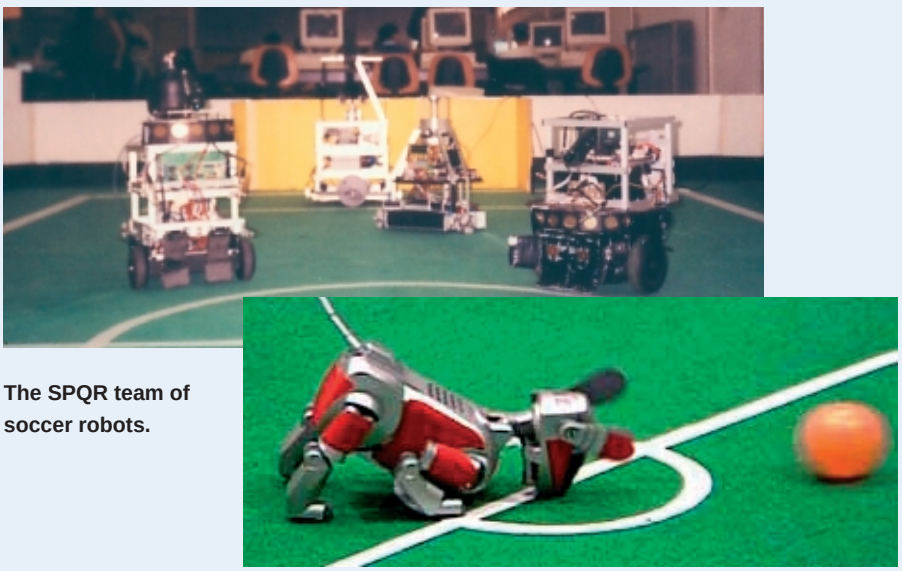

The SPQR team of soccer robots.

Links:
http://www.robocup.org
http://www.dis.uniroma1.it/~spqr/

Please contact:
Luca Iocchi,
University of Rome "La Sapienza", Italy
Tel: +39 06 8841947
E-mail: iocchi@dis.uniroma1.it

3APL - A Programming Language for Cognitive Agents

by Mehdi Dastani, Frank Dignum and John-Jules Meyer

The 3APL project (An Abstract Agent Programming Language) aims at developing a programming language to implement cognitive agents and the high level control of cognitive robots.

Many complex applications, such as electronic auctions, electronic markets, or autonomous robots, can be analyzed more intuitively and directly in terms of interacting cognitive agents using high-level concepts such as beliefs, goals, desires, intentions, obligations, actions, plans, and communication. Various tools, specification languages, and architectures are proposed to analyze, specify, and design such cognitive agent systems. However, there has been considerably less attention to developing programming languages for cognitive agents that provide programming constructs to implement high-level concepts directly. This is the objective of the 3APL project.

The 3APL project is carried out at Utrecht University (Institute of Information and Computing Sciences, Intelligent Systems Group) and is partially supported by the Dutch Research Council NWO. The work is closely related to other ongoing research projects such as epistemic and dynamic logics within our group, and AgentSpeak(L) and Golog/Congolog projects outside our group. Although the aim of the AgentSpeak(L) project, carried out at the Australian Artificial Intelligence Institute, was similar to that of 3APL, only a limited set of cognitive concepts can be implemented by their proposed programming language. The aim of Golog/Congolog, under way at Toronto University and Aachen University of Technology, is to develop a programming language to implement the high level control of cognitive robots in dynamic and unpredictable environments. While 3APL provides programming constructs to implement a large set of cognitive concepts, Golog/Congolog concentrates on programming constructs, such as sensing and planning, which are essential for dynamic and unpredictable environments.

The 3APL project has developed a programming language for implementing cognitive agents with beliefs, observations, actions, goals, communication, and reasoning rules. In particular, agents' observations and beliefs can be implemented in 3APL by a subset of first-order predicate language (prolog-like facts and rules). The actions are implemented as triples consisting of action name together with pre- and post-conditions. The pre- and post-conditions of actions are belief formulae indicating the condition under which the action can be performed and the effect of the action after it is performed, respectively. The goals that can be implemented in 3APL are procedural or to-do goals, which can be implemented by expressions of an imperative language. These expressions are formed by applying constructs such as sequence, test, conditional choice, and recursion to actions and belief formulae. Communication can be implemented by the pre-defined send-message action. Finally, reasoning rules can be used to implement the generation of goals, the revision of actions that are blocked, the revision of goals that are not achievable, the optimization of goals, etc.

We have also developed an interpreter that executes 3APL programs. The interpreter implements the so-called deliberation process that involves many activities such as applying reasoning rules, selecting goals to be achieved and actions to be executed. The interpreter is implemented in JAVA. A JAVA class, called 3APL Class, has been developed. This class has private attributes that represent cognitive concepts such as beliefs, goals, basic actions, and practical reasoning rules. The 3APL class has two main methods: one method for the creation of the 3APL object through which the programmed cognitive concepts are parsed and represented by the corresponding private attributes, and one method for the actual execution of the object. The latter method corresponds to the deliberation cycle of 3APL. The architecture of the cognitive agents that can be programmed by 3APL is illustrated in Figure 1.

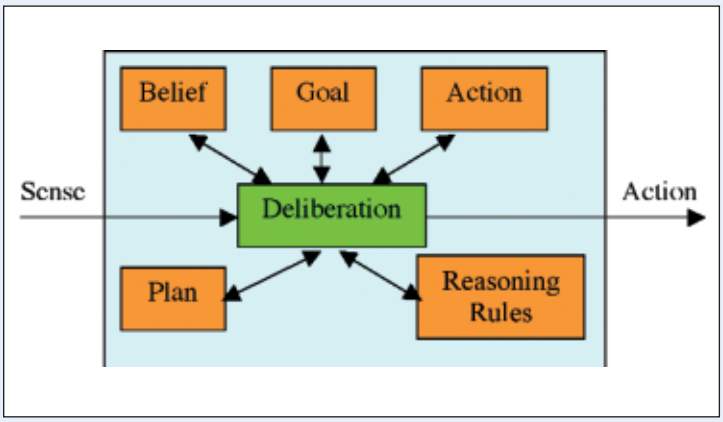

The 3APL Agent Architecture.

Different applications require different deliberation processes, ie there is no single universal deliberation process. Therefore, the deliberation process of cognitive agents should be controlled by the programmer. This implies that the deliberation process itself should be programmable. We have thus designed two programming languages: an object-level and a meta-level language. The

object-level language has been recently extended with programming constructs to implement declarative (to-be) goals and plans. The meta-level language includes terms that refer to object-level entities such as goals, actions, plans, and rules, and imperative programming constructs to implement how to process the object-level entities. The current state of the 3APL project provides the formal specification of the syntax and semantics of the two programming languages. The semantics of these languages are defined in terms of transition systems.

We have used 3APL to implement the high level control of moving robots. This is achieved by extending the object-level language with physical actions. One problem with physical actions is that their post-conditions are not known beforehand. We have coupled the 3APL interpreter to the control software (ARIA) of activmedia robots. This coupling translates 3APL basic actions to robot actions, and uses the robot's sensor information to set the post-conditions of the physical actions. Several future activities are planned, including the extension of the interpreter to execute the declarative goals and the meta-level programs, and setting communication between agents according to the FIPA (Foundation for Intelligent Physical Agents) standards. Moreover, we are implementing an agent platform that can be installed on various machines. Agents from different platforms can communicate. Various materials including the current interpreter, some examples of 3APL programs, and various publications are available on the project web page.

Link
Project web page:
http://www.cs.uu.nl/3apl/

Please contact:
Mehdi Dastani, University of Utrecht,
The Netherlands
Tel: +31 30 253 3599
E-mail: mehdi@cs.uu.nl
http://www.cs.uu.nl/~mehdi

Environment Representations for Cognitive Robot Companions

by Ben Kröse

World representations for intelligent autonomous robots are studied at the University of Amsterdam. An omni-directional vision system is used in combination with appearance modelling for robot localization and navigation. The goal is to develop a robot interface between human and intelligent environment. Also robot learning of concepts by interaction with humans is investigated.

An increasing amount of research is spent on robots that operate in natural, non-industrial environments and interact with humans such as household robots, sociable partners, education robots or personal assistants. Since the natural environment is highly unstructured and possibly unknown, these robots cope with the problem of learning an internal model of this environment. Furthermore, since these robots are to interact with humans in a natural way, the representation should correspond with human concepts of space.

At the Informatics Institute of the University of Amsterdam several research projects are focussed on world representations for intelligent autonomous systems. Such systems perceive the environment, reason about it, and act in it. An internal representation is needed for planning and navigation. For example, if a household robot is instructed to vacuum the kitchen, it needs to use the internal model to plan a path from its current location to the kitchen.

Conventional industrial mobile robots use geometric models, like polygons or occupancy grids, which fit very well with the range sensors traditionally used on these types of robots. Although adequate for industrial environments, which are generally simple and have a regular structure, more powerful models are needed for natural environments like houses, of which the structure may be complex and not easy to describe with simple geometric primitives. In many cases also other characteristics of the environment such as colour or texture need to be modelled, which may be a more powerful cue for robot location than geometric properties. Luckily today's developments in computational power and sensor technology open a way for the use of vision systems on mobile robots. We investigated whether a vision system can also be used for robot localization. Can the robot 'see' where it is?

In our studies we use an omni-directional vision system, which is able to capture images of the full surroundings at video rate. The issue that arises is which type of internal representation is best suited for the (navigation) task of the robot and which best corresponds to our visual sensor. One approach is to use a full CAD model, which combines the geometry of the environment with intensity and colour. However, such models are difficult to derive from visual information. We took an alternative approach where we do not explicitly model the environment, but made an implicit environment representation by modelling the relation between the observed image and the robot location. This is called 'appearance' modelling. Such a model can be learned more easily than a CAD model, and we showed that it is well suited for the navigation task.

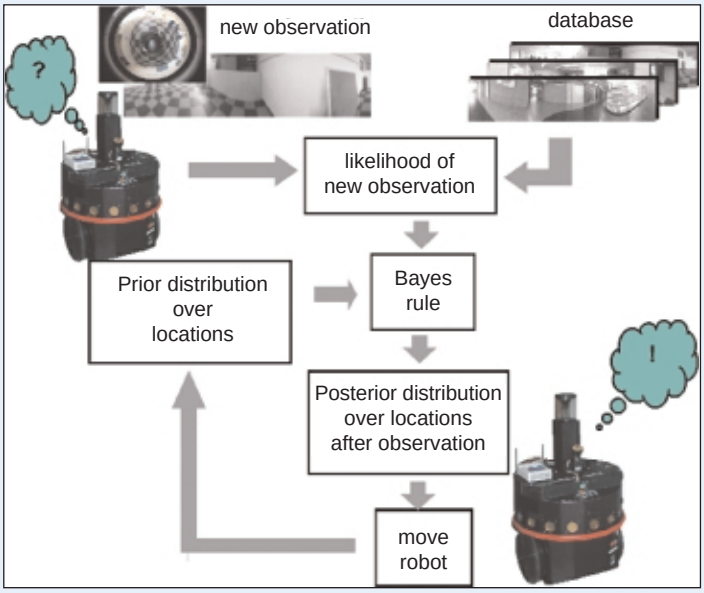

Figure 1: The Markovian robot localization. To calculate the likelihood of a new observation, a kernel estimator is used on the images from the database. The distribution of the robot location is implemented as a Monte Carlo method.

In appearance modelling a large set of images taken at different locations is stored. For our office environments we typically use about 1000 images. A problem with modelling in the image space is the high dimensionality of the data. Our omni-directional images are typically 64 x 512 pixels, so an image is a vector in a 32K-dimensional space. We therefore first reduce the dimensionality of the images with a Principal Component Analysis (PCA), and store the linear features. For robot localization we use a probabilistic framework. The current observation of the robot is compared with the database and the likelihood of the observation is determined. Since the likelihood is conditioned on the location of the robot, we can determine the posterior probability distribution of the robot position by using Bayes rule and the prior distribution of the robot positions, which is estimated from the previous estimate. The localization procedure is schematically depicted in Figure 1. An important part of the research project is to find a computationally efficient way to represent the probability distributions. We used a special variant of a Monte Carlo 'particle' filter.

This method was developed within a large project funded by the Japanese MITI, in which we collaborated with various Japanese research laboratories, and was successfully demonstrated at various fairs and exhibitions. Currently the research is continued in a European project on Ambient Intelligence, in cooperation with Philips, Epictoid (a CWI spin-off) and many other European partners. Here the goal is to make a robot, which serves as interface between the human and the intelligent environment. Such a robot must be able to navigate in the home environment and to interact with humans in a natural way by speech understanding, speech synthesis and an expressive face. In this project we use the appearance modelling approach for localization and navigation.

Although this type of environment modelling is a step further on the way to a more natural representation, we do not have yet the concepts, which can be used to communicate with humans. In a project funded by the Netherlands Organization for Scientific Research NWO we investigate whether the robot can learn 'concepts' by interaction with humans. The objective of the project is to develop learning methods, which are able to form 'state-action' concepts, such as: 'go to the door, then left, then after the second fire extinguisher right', or 'perceptual' concepts, such as: 'corridor', 'hallway' or 'window'. Again we use the omni-directional vision system, and we developed fast learning methods to enable real-time learning. The first experiments show promising results, and we hope to have a robot, which learns elementary concepts in a short training time. The challenge is to show that an embodied intelligent agent is able to learn concepts by navigating in real environments and interacting with humans.

Please contact:
Ben Kröse
University of Amsterdam, The Netherlands
Tel: +31 20 525 7520/7463
E-mail: krose@science.uva.nl
http://www.science.uva.nl/~krose/

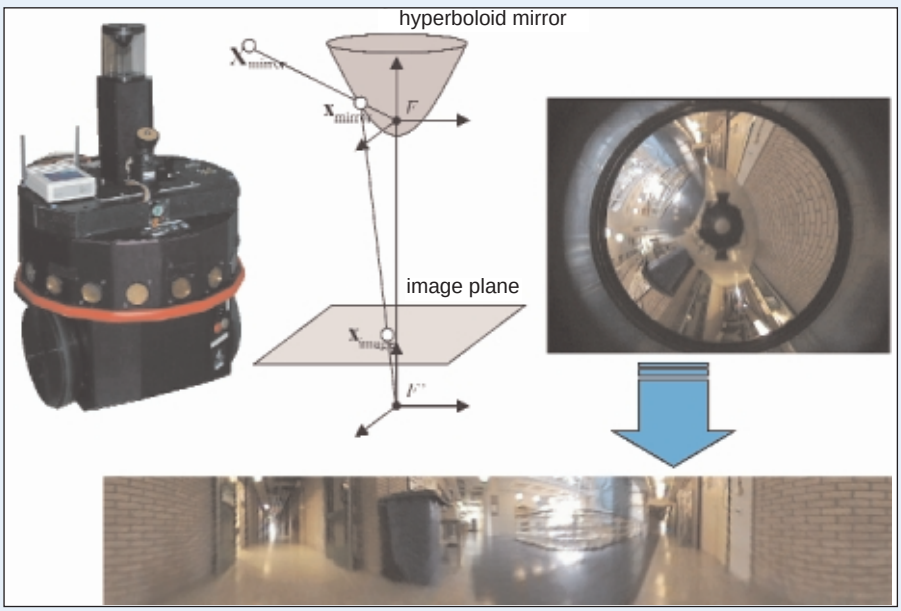

Figure 2: Left: the mobile robot with the omnidirectional vision system. The properties of the hyperboloid mirror enable an omnidirectional image (upper right) to be reprojected easily onto other surfaces, for example onto a virtual cylinder, resulting in a panoramic image (bottom).

Towards Tele-Presence: Combining Ambient Awareness, Augmented Reality and Autonomous Robots

by Pieter Jonker and Jurjen Caarls

A setup for a mobile user employing Augmented Reality for indoor and outdoor applications has been developed at Delft University of Technology. To determine the user's position and orientation in the world, it uses a cognition system similar to that used by our autonomous soccer playing robots. In order to let robot and man cooperate, a knowledge system shared by man and robot is under development, in which both can express their understanding of the 'world' as they perceive it. Applications include rescue operations in hazardous situations.

The proliferation of digital communication systems, infrastructures, and services will eventually lead to a Personal Communicating Digital Assistant (PCDA): a hybrid between mobile phones and Personal Digital Assistants with ubiquitous computing capabilities. A PCDA possesses Ambient Awareness: the ability to acquire, process, and act upon application specific contextual information, taking the current user preferences and state of mind into account. A major issue of Ambient Awareness is positioning: the device must know where it is. For outdoor navigation DGPS is a good candidate to detect one's position. Indoors it is easier to set up a network of beacons. Radio beacons can realize good accuracies. An alternative is to use a camera and a model of the world. From features in the image a position can be calculated and tracked, in the cm range. Such features (visual beacons) can be nameplates on doors, or dot patterns fixed at known positions in a building.

A PCDA also employs Augmented Reality: a context aware application (Figure 1) which merges virtual audio-visual information with the real world using a see-through display (Figure 2). Outdoor DGPS or indoor a beacon system captures the user's rough position. A camera captures the user's environment, which, combined with gyroscopes, accelerometers, and compass, makes the PCDA fully aware of the user's absolute position and orientation with such an accuracy that virtual objects can be projected over the user's real world, without causing motion sickness. Camera images are sent to the backbone and matched to a 3D description of the environment derived from a GIS database of the environment, to determine the user's position.

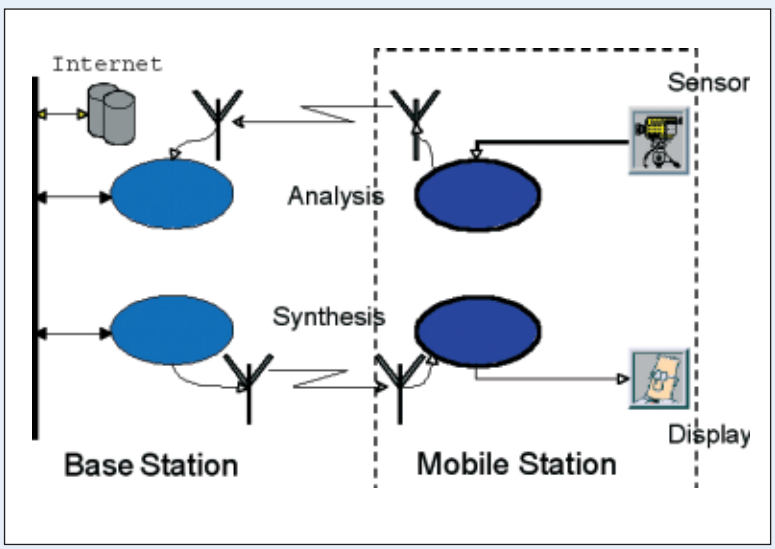

Figure 1: A PCDA using network services.

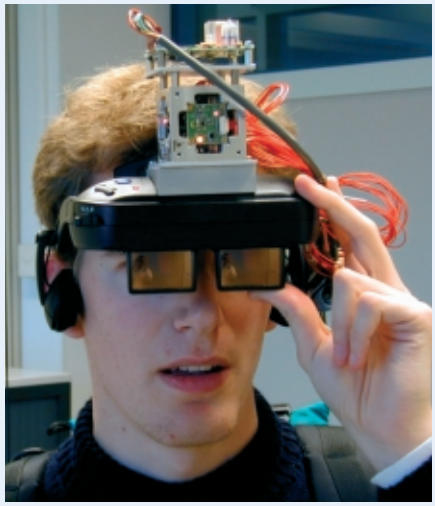

Figure 2: Prototype Augmented Reality headset.

The fusion of the sensors, each with their own accuracy, update rate and latency, is based on Kalman filtering. With our combined inertia and visual tracking system we are able to track head rotation and position with an update rate of 10 ms with an accuracy for the rotation of about 2 degrees, whereas head position accuracy is in the order of a few cm at a visual cue distance of less than 3 m. In our experimental set-up a background process, once initiated by the DGPS system, continuously looks in the image for visual cues and - when found - tries to track them, to continuously adjust the drift in the inertial sensor system (Figure 3). Our set-up is a context aware system that can run on future generation PCDAs for professional use. Low-cost versions can be derived with camera and some position sensing devices, but without Augmented Reality. This lowers the requirements for the positioning accuracy and update rate drastically. In such a consumer version the camera can be used to realize the awareness of the user's position (where is he/she?), the user's attention (what is the user looking at, pointing at, maybe thinking about?),

and the user's wishes (what is the problem, what information is needed?).

The annual world championship Robot Soccer (see Figure 4) is organized by the international RoboCup organization (www.robocup.org). One of the most important features of the robots is their cognition system. The determination of the robot's own position is based on its odometry augmented with its visual system that, based on field lines and cues such as the goal colours, can continuously calibrate the odometry. In about ten years walking robots (the Humanoid League) will take the lead.

In a second domain of RoboCup, in the Rescue Simulation League, the challenge is to rescue as many civilians and buildings as possible after a disaster in a city, such as the Kobe earthquake (Fig. 5). This agent based simulation contest is a preparation for the real Rescue Robot troops. One can think of, in analogy with a dog brigade in a police force, a troop of robots under command of the fire brigade that helps to discover and save injured, to discover and extinguish fires and help to set-up communication and information spots.

Of crucial importance is, how the robots in the simulation build up their vision on the world (E-Semble, see www.e-semble.com) and how they can share information using limited communication channels (Figure 6). This is an extended problem of the real robots playing soccer, with the difference that in a simulated world only symbolic knowledge need to be treated, whereas in a real world also perceptions play a role. As with humans, perceptions of distance, size and shape are inaccurate, because the goal is not to obtain an accurate scene description and measurement, but rather acquisition of task specific information for object recognition, obstacle avoidance and object handling.

We have realised an Augmented Reality system, a team of autonomous soccer playing physical robots and are involved in the simulation of disasters. If Autonomous Robots, Ambient Awareness and Augmented Reality devices become more mature, a quite natural step would be to let the humans

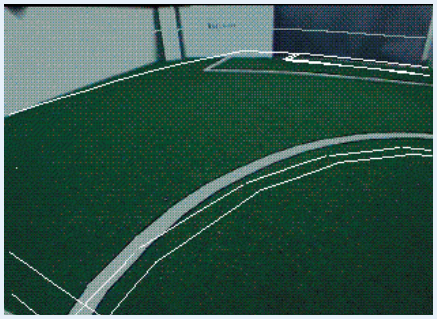

Figure 3: Matching of visual cues with a world model to determine a device's position and orientation.

Figure 4: Autonomous Soccer playing Robots using cognition systems for Ambient Awareness.

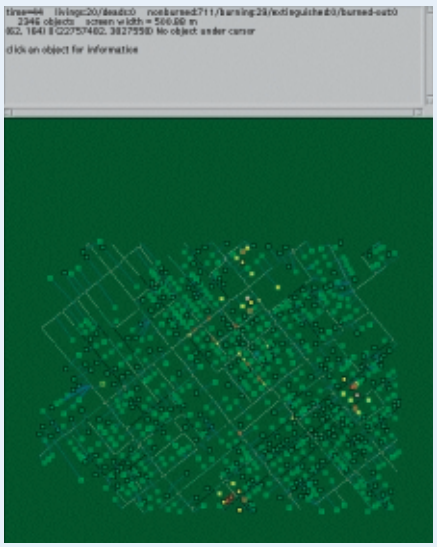

Figure 5: RoboCup Rescue Simulations: A disaster management competition.

Figure 6: A local view on a disaster.

work side to side with robots, eg for search and rescue. The humans, fitted with Augmented Reality PCDAs, could be made to perceive what the robots perceive on far more dangerous grounds, ie, tele-presence. This, however, requires a mutual knowledge system for symbolic knowledge (facts) as well as perceptual knowledge. The symbolic knowledge must contain data descriptions of fixed and dynamic objects, their attributes and the relations between the objects. The perceptions on objects must include more vague descriptions, eg, based on support vector data descriptions of the objects. As the system covers a dynamically changing environment, it must be able to learn and forget symbolic knowledge as well as perceptions. Crucial is that the system must be able to ground perceptions to symbols (ie label it and relate it to facts). For example a vague black blob, perceived by some robots can be labelled 'door' by a human, whereafter the robots can use this fact in their world model. Although at this moment there exist many components of systems that maintain a world, their integration for tele-presence is ongoing.

This work is performed in cooperation with Delft University of Technology (faculty ITS; Augmented Reality), the Telematics Institute Enschede (Ambient Awareness), the University of Amsterdam (faculty of Science; Robot Soccer) and E-Semble (Rescue Simulation).

Please contact:
Pieter Jonker or Jurjen Caarls
Delft University of Technology,
The Netherlands
Tel: +31 15 278 3763/ 3727
E-mail: {pieter, jurjen}@ph.tn.tudelft.nl

The RoboCare Project: Multi-Agent Systems for the Care of the Elderly

by Amedeo Cesta and Federico Pecora

The RoboCare project focuses on the development of distributed systems in which software and robotic agents contribute to the common goal of generating active services in environments in which humans may need assistance and guidance, such as health care facilities.

Thanks to recent technological advances, there are now many potential applications for robotics including multi-agent systems. The aim of the RoboCare project is to study issues and challenges involved in the design of systems for the care of the elderly that adopt both fixed and mobile heterogeneous agents. These agents can be robots, intelligent sensors or possibly even humans. RoboCare shares some aspects with other projects for the assistance of elderly people, such as Pearle, the mobile robotic assistant for the elderly (M. Pollack et. al.) and the Assisted Cognition Project (H. Kautz et. al.), but addresses the particular goal of creating a heterogeneous multi-agent environment for generating user services.

The ambitious goal of providing support for the elderly with the use of intelligent systems is not merely aimed at automating certain tasks. The philosophy underlying this application of cognitive technology is to enhance the quality of care for the elderly, by employing an infrastructure of expert systems which cooperate in order to provide services. This requires, on one hand, the design of intelligent agents endowed with advanced learning capabilities, capable of complex symbolic reasoning tasks and high levels of interaction with humans. But in a world in which pervasive computing is the name of the game, the integration of intelligent capabilities in aiding tools is no longer enough. Tomorrow's health-care institution will be equipped with intelligent systems which are capable of interacting with humans; RoboCare will address the issue of enabling these systems to work together.

The implementation of collaborative behaviours for intelligent systems is motivated by the necessity of providing constant care and support for people with disabilities. When it comes to caring for the elderly, intelligent systems aim at supporting them in all aspects of daily life, providing decision making, reminding and warning functionality. Such complex tasks necessarily involve teams of agents (be they robots, intelligent sensors or physicians). For instance, Alzheimer patients would benefit from a system which reminds them about certain sequences of actions to be performed (such as switching off the gas after turning off the stove and wearing glasses before wandering out of the bedroom). In its simplest form, such a system would be made up of a series of agents, like a gas monitor and a mobile robot capable of advising the assisted person. All the actors in the system would clearly be capable of carrying out individual reasoning, but would also need to collectively reason about the situations which can occur.

Creating such tightly coupled intelligent systems presents an important challenge for the AI and robotics communities. On one hand, researchers have been successful in creating teams of robots which are capable of cooperating in tasks such as foraging, unknown environment exploration and simplified forms of soccer playing. On the other hand, the AI community has advanced in the study of planning and scheduling, and has developed mature technology for automated problem solving. The realization of a complete system for the aid of the elderly to be employed in a complex, real world environment such as a health-care institution or a home requires a tight integration of these technologies. The

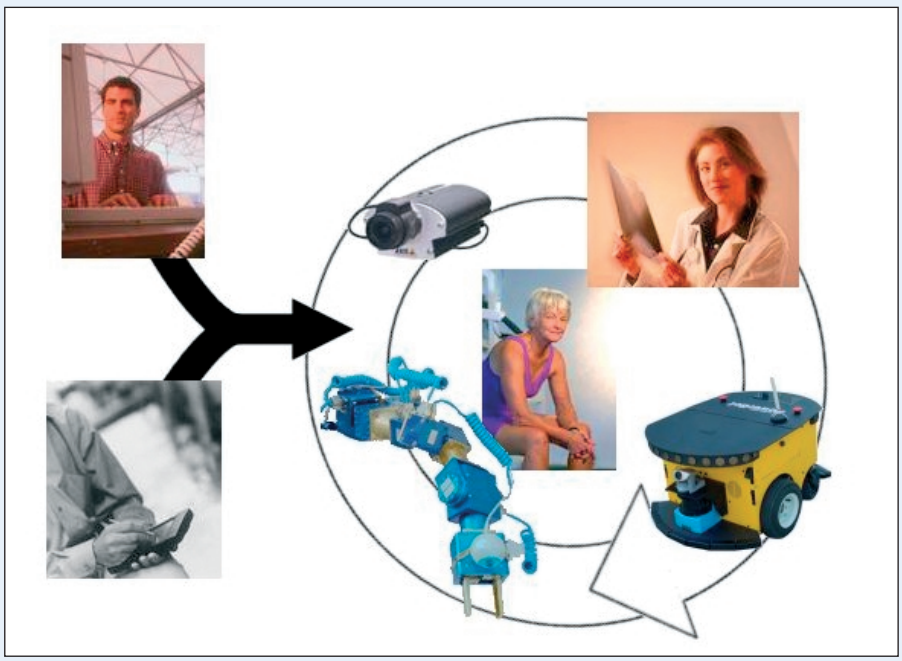

Overall system structure: operator level system control and distributed environment.

system must provide more than a loose coordination between components, in other words, it must enable a synergetic cooperation among the agents. Thus, a supervision framework which maintains a global view of the system and provides control functionalities for human operators must be provided. On the other hand, in order to ensure robustness and reliability, the autonomous agents must be endowed with well-defined cooperative behaviours.

However, the task of providing monitoring support is but a small aspect of the functionalities that such a system must implement. Caring for elderly people is often all about routine, such as scheduling walks and daily social events. As a consequence, the system should be capable of planning and scheduling for such needs as well as interpreting them.

A great deal of resources are being invested in the investigation of user-acceptance of the technology, in terms of the psychological impact of robotic agents providing services for the care of the elderly. Since the level of social interaction which can be obtained between robotic agents and human beings is related to familiarity, we are also investigating the use of Aibo™ robots for monitoring tasks.

Let us conclude by saying that the issues which need to be addressed in the context of complex systems for the care of the elderly are communication, knowledge representation, human-machine interaction, learning, collective and individual symbolic reasoning, only to name a few. RoboCare is made possible by the cooperation of a number of participating units, which include leading Italian universities and various departments of the National Research Council (CNR). Thanks to the numerous units involved, the RoboCare project can count on experts from many fields of cognitive science and technology.

RoboCare is funded by MIUR (the Italian Ministry for Research and Education) and is coordinated by the Institute for Cognitive Science and Technology of CNR.

Please contact:
Amedeo Cesta, ISTC-CNR
E-mail: cesta@ip.rm.cnr.it

Robots and Girls - A Promising Alliance

by Monika Müllerburg and Ulrike Petersen

Education and training are important requirements for the future abilities of a society. This especially concerns key areas of technology such as information science and the engineering disciplines. We notice a decreasing interest from students in technical professions, and as a consequence there is a remarkable lack of junior scientists. If we want to convince more young people to choose technical professions, we need to change the teaching program. Girls and women in particular need to be addressed.

Germany is falling behind its European neighbours in terms of the participation of women in science. There is a lack of female students, particularly in technical disciplines. In general, women avoid studying topics like engineering, electronics and physics; in 2000, the quota of female students was only about 17%. Women are more attracted by interdisciplinary topics such as architecture, multimedia and environmental engineering.

It is evident that engineering and computer science are key professions for the European economy. It is necessary that women should participate in the creation of future technologies, and our society must seriously consider the question of whether it can afford not to utilise the potential talents of women.

To increase the number of women in these fields in the long run we must first address girls. Education and training must inspire girls to take up technical subjects. Our experience shows that edutainment robotics meets with great interest, and that learning by doing is a promising way of raising the appeal of technical subjects, especially to girls and women.

This is one of the main aims of the project 'Roberta - Girls Discover Robots'. The name 'Roberta' was chosen as an association for a female robot. The project will establish courses of instruction to design, construct and program robots. Particular attention will be paid to gender aspects. The challenge is to define topics which meet the special interest of girls. For example, since girls are more interested in the solution of

Subject	Total	Women	% Women	Men	% Men
Electrical engineering	11.350	1.018	9.0%	10.332	91,0%
Computer Science	27.157	4.958	18,3&	22.199	81,7%
Machine-building	20.905	3.860	18,5%	17.045	81,5%
Total/average	59.412	9.836	16,56 %	49.576	83,44 %

Beginners at German universities in technical courses of studies (year: 2000). Source: Federal Statistical Office Germany, Wiesbaden 2001.

No sweet without sweat.

environmental problems, we will deal with subjects that simulate natural phenomena. We know by experience that these topics will interest boys as well.

The 'Roberta' team will develop didactic measures, teaching and studying papers and make them available for teachers and tutors. Web-based background information and electronic documents will be included. In addition to traditional teaching, the girls will do practical work with robotic kits. They will learn to design and to construct robots, to write programs and to combine these activities within a system development process.

Both the papers and the courses will be evaluated, which will help to improve the results. Special attention will be put on female views and approaches to stress the gender quality of the project.

Regional centres will be established to support the distribution of the results beyond the end of the project. Teachers are invited to join these and to exchange their experiences. The aim is to run a national network, which will ensure the further progress of the development of gender criteria.

'Roberta' is funded by the German Federal Ministry of Education and Research. Partners include universities, a technical museum, a school board, a grammar school, the Centre for Women in Information Society and Technology and the LEGO Educational Division.

Project 'Roberta' is a part of a broader initiative running at the Fraunhofer Institute for Autonomous Intelligent Systems (AIS). Edutainment robotics is one of the business fields at the AIS, where robot construction kits are developed. The idea is to use robots and robotic construction kits to teach engineering know-how.

The 'Roberta' team is interested in cooperation with similar initiatives in other European countries and would like to exchange experiences and define joint activities in the field of education based on robots.

Links:
www.ais.fraunhofer.de/ROCK/roberta
www.ais.fraunhofer.de/

Please contact:
Monika Müllerburg, FhG-AIS
Tel: +49 2241 14 2428
E-mail: Monika.Muellerburg@ais.fraunhofer.de

Man and Machine

by Brian R. Duffy

The primary research challenges of the Anthropos project are to understand the relationship and to establish a bond between man and machine. Can the illusion of life and intelligence emerge through simply engaging people in social interaction with a robot?

AI research to date has led to a wealth of strategies for solving local problems with varying degrees of success. The next step has been to integrate these tools into whole systems, such as robotics, to realise a system that displays a breadth of 'intelligent' behaviour. This integration proved to be nontrivial and highlighted issues with classical AI such as its strong representation approaches and often narrow take on intelligence. In looking to extend beyond traditional paradigms in developing an artificial form of 'intelligence', the field of social robotics has recently emerged.

An impressive display of human intelligence is in our ability to adapt in social scenarios. In fact, social intelligence has been theorised as the mechanism from which the capability for the complexity of human languages evolved. Thus, one of the purposes of Media Lab Europe's research on socially-capable robots is to investigate human social intelligence and behaviour. From a human-machine interaction perspective, a socially-capable robot facilitates our access to the digital world through intuitive social mechanisms to improve or provide alternative approaches in education, information dissemination, and our future daily interactions with machines. Once

Photo courtesy of Brent Jones

Media Lab Europe's 'JoeRobot' at the Flutterfugue performance with SmartLab and NYU CATLab in London 2002.

domestic robots become more than just a washing machine in the home, our social interaction with them becomes inevitable. The issue is how do we realise such a socially capable robot?

In order to begin answering these questions, the Anthropos Project at Media Lab Europe seeks to decompose the interaction issues between man and machine. The current research areas are:

- Balancing function and form for social robots. Anthropos and JoeRobot (Figure 1) are prototypes built to explore the development of socially capable robots. Key to the notion of expandability and rapid prototyping, a modular nervous systems strategy employs standardised interface protocols (Firewire and USB) for actuator and perceptor components. Research on integrating a socially capable robot into performance spaces has demonstrated the power of the form as an interface to the digital information domain (figure 1). People's willingness to engage with a machine that judiciously employs anthropomorphic features we are familiar with in social contexts facilitates man-machine interaction.
- Strength and degree of minimal expression and communication. The Emotion Robots work is a series of experiments to investigate how minimal the set of humanlike features can be for a social robot. Data illustrates people's propensity to attribute such concepts as emotions and intelligence to machines performing computationally simple behaviours.
- The seamless integration of real physical worlds and digital information space. The Agent Chameleons project strives to develop digital 'minds' that can seamlessly migrate, mutate and evolve on their journey between and within physical and digital information spaces. This challenges the traditional boundaries between the physical and the virtual through the empowerment of mobile agents. Three key attributes mutation, migration and evolution underpin the Agent Chameleons concept where digital personal assistants are developed that opportunistically migrate and choose a 'body' (whether a robot, an avatar in virtual reality, an animated character on a PDA, or a web agent) to facilitate its intentions.

The Anthropos Project draws this work together to look at developing a different perspective on what an artificially intelligent entity could become. Machines have intrinsic properties that are often seen as hindrances when the reference is either humans or other biological entities. The objective is to embrace those aspects that are constructive and integrate these with a machine's inherent advantages, ie being a machine. While arguments have prevailed for many years over the nature of intelligence and whether it can be realised in a machine, this work aims to demonstrate the power of perceived intelligence and people's willingness to interpret a social robots interactions according to human-like social references. The key issue becomes a balance between function and form.

Links:
Anthropos Project at Media Lab Europe:
http://anthropos.mle.ie

Please contact:
Brian R. Duffy
Media Lab Europe, Dublin, Ireland
Tel: +353 1 474 2823
E-mail: brd@media.mit.edu

Enhancing Non Player Characters in Computer Games using Psychological Models

by Brian Mac Namee and Pádraig Cunningham

In the ever-growing research field of computer games, one of the main research areas is the creation of believable non-player characters. A part of the TCD Game AI Project, is furthering this research by creating a connectionist system for the simulation of social interactions amongst non-player characters based on psychological models of personalities, moods and relationships.

In modern computer games the level of graphical realism is so advanced that players can be led to believe that games are set within realistic game worlds. However, this illusion is often shattered as soon as the player begins interacting with computer controlled non-player characters (NPCs). The TCD Game AI Project at Trinity College Dublin is applying artificial intelligence and machine learning techniques to the problem of creating believable NPCs.

One of the areas in which NPCs are lacking is the kind of interactions in which they engage, both with other NPCs and players. An exception to this rule, The Sims (thesims.ea.com), has shown that games which make management of character interactions a focus of game play can be highly successful.

The µ-SIC system (which is a part of an agent architecture designed specifically for the creation of NPCs) uses quantitative psychological models to capture NPCs' personalities, moods and relationships. The values of these models are used as inputs to an artificial neural network (ANN), which drives characters' social behaviour causing them to perform interactions such as flirting, joking, and chatting

Modelling Personality, Mood & Relationships

In order to capture the important aspects of NPCs' personae the following quantitative psychological models are used:

- Personality Model: To model NPCs' personalities Eysenck's classification model has been chosen. This plots personality across two orthogonal axes, introversion-extroversion and neuroticism-stability, allowing the creation of characters with personality types, such as aggressive, sociable and moody.
- Mood Model: To simulate a character's mood a model from Lang is used. Agents' moods are measured according to valance and arousal, where valance refers to whether the mood is positive or negative, and arousal refers to the intensity of the mood.
- Relationship Model: Agents' relationships are simulated using a model with its psychological basis in work undertaken by Wish. Typically this models plots the relationship between two characters across four axes: the amount that a particular character likes another character, physical attraction, dominance or submissiveness and intimacy. To facilitate conversation, this model has been augmented with a value indicating how interested one character is in another. A high interest rating indicates that characters share a number of common subjects of interest, and are thus more likely to converse.

Implementing the µ-SIC System

An ANN is used to transfer a particular set of a values for the models just described to a particular interaction for an NPC to engage in at a particular moment. The structure of the ANN used by µ-SIC (a multi-layer perceptron with a single hidden layer) is shown in figure 1. The network has been trained using an artificially created data set based on interactions predicted by a group of researchers.

In order to reduce the storage requirements of the µ-SIC system, only one copy of it is stored within a game engine. Whenever NPCs are free to begin an interaction they query this master copy of the µ-SIC system and so it can be considered an oracle that advises NPCs on how to behave.

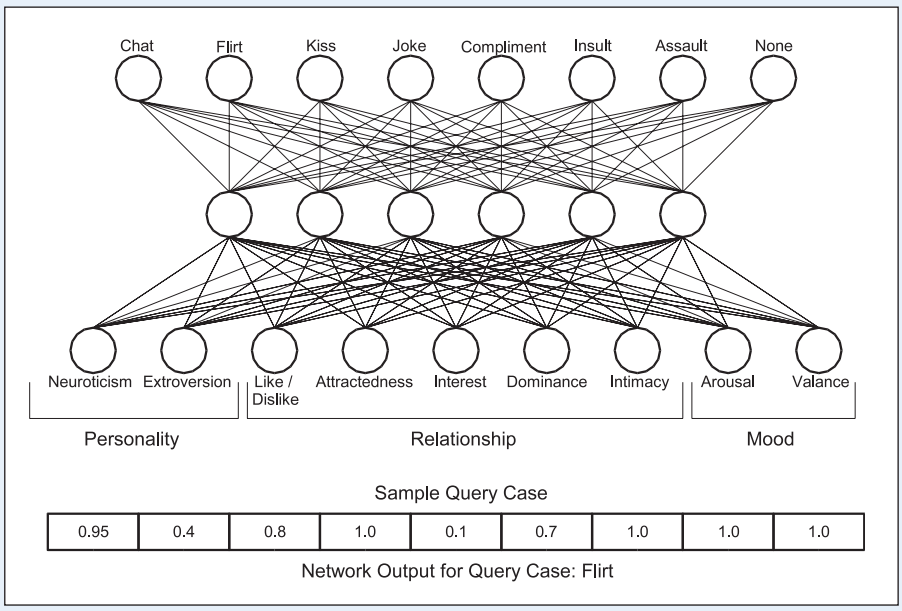

Figure 1: The structure of the ANN used within µ-SIC.

Figure 2:
A screenshot of the simulation example demonstrating the µ-SIC system.

Simulation Example Evaluation

To evaluate the performance of the µ-SIC system a simulation example (graphically represented in a very simple, cartoon style) which takes place in a small town has been constructed.

The simulation is populated by a range of NPCs, each of which has a personality defined using the Eysenck model. The system runs in real-time with NPCs moving between different locations performing day-to-day tasks such as going to work and visiting restaurants and bars. At various times during the simulation NPCs are free to perform interactions chosen by the µ-SIC system, which in turn cause their moods and relationships to adjust. A screenshot of the simulation running is shown in figure 2.

Conclusions

The µ-SIC system can be used as part of a larger agent architecture to allow computer game NPCs perform social interactions. These interactions are based on characters' personalities, moods and relationships which are captured using quantitative psychological models. Interactions are chosen using an ANN which has been trained to determine an appropriate interaction based on particular values of the models used. Within a simulation example created characters perform the full range of interactions which are always consistent with their personalities and current mood and relationships with other characters. To improve the system extra inputs will be added to the ANN. These include the last interaction performed by the character and values indicating NPCs' current locations which should block certain interactions. The system is also being integrated into the Trinity College Image Synthesis Group's ALOHA system, which performs sophisticated animation of virtual humans.

Link:
http://www.cs.tcd.ie/research_groups/aig/GameAI/

Please contact:
Pádraig Cunningham, Trinity College Dublin
Tel: +353 1 608 1136
E-mail: Padraig.Cunningham@cs.tcd.ie
http://www.cs.tcd.ie/Padraig.Cunningham/

Understanding and Interpreting the Activities of Experts: Towards a Cognitive Vision Methodology

by Stelios Orphanoudakis, Antonis Argyros and Markus Vincze

The ActIPret project, funded by the European IST programme, aims to build advanced vision systems, able to recognize and interpret the activities of experts in the context of a cognitive vision framework.

One of the fundamental abilities possessed by humans is that of acquiring skills through observation. Teaching by demonstration is consequently a powerful way to provide training. Despite its potential, this type of teaching is not always possible because of distance and time barriers; experts can only teach small groups of trainees, at a certain location and for a limited time period. These barriers could be removed if we could realise a computational vision system that is capable of understanding, interpreting, storing and indexing the activities of experts. The combination of such a system with recent advances in the field of virtual and augmented reality (VR/AR) could be used to effectively search, retrieve and realistically reproduce the activities of experts anywhere and anytime.

Humans have a remarkable ability to visually interpret the activities of other humans, transform these interpretations into knowledge and subsequently exploit this knowledge in acquiring related skills. Teaching by demonstration therefore constitutes a powerful training technique. Currently, teaching by demonstration entails an expert demonstrating her/his expertise to small groups of trainees on specific and limited occasions. Time and distance barriers hinder the observation of experts in action. Recording experts' activities on video provides a partial solution to this problem and allows repeated viewing but is subject to other important limitations. Although video records a dynamic sequence of events, in some sense it is quite a 'static' source of information. The fixed viewpoint restricts visibility and may lead to ambiguous interpretation. Moreover, activities cannot be indexed and effectively searched, as is the case with information in the form of a manual or user's guide.

With recent developments in virtual and augmented reality (VR/AR), it is now possible to produce high-quality representations of a reconstructed scene and a realistic replay of activities therein. Such capabilities could prove invaluable in developing tools for teaching by observation, provided that the recognition and interpretation of the activities of an expert is also possible. The coupling of these capabilities could result in the removal of most of the important barriers in teaching through observation. The experts' demonstration can be replayed anywhere, anytime and from any viewpoint. Moreover, the activities can be indexed efficiently and effectively, and retrieved by the trainee based on her/his needs.

The overall objective of ActIPret is to develop a cognitive vision methodology that permits the recording and interpretation of the activities of people handling tools. Interpreted activities can be stored in an activity plan that can be referenced later by the user. The activity plan is an indexed manual in the form of 3D reconstructed scenes, which can be replayed at any time and location to many users using VR/AR equipment. Research and development is focused on the active observation and interpretation of the activities, on the extraction of the essential activities and their functional dependence, and on organising them into constituent behaviour elements. The approach is active in the sense that the system seeks to obtain views that facilitate the interpretation of the observed activities. Moreover, task and context knowledge is exploited as a means of constraining interpretation. Robust perception and interpretation of activities is the key to capturing the essential information, allowing the reproduction of task sequences from easy-to-understand representations and providing a user-friendly tool for the trainee.

The figure illustrates the envisioned scenarios of use of the ActIPret system. During recording, the expert's activities are observed and an activity plan is obtained. During replay, the trainee/user searches for specific activities of interest using a conceptual language. The user is then able to choose between two options:

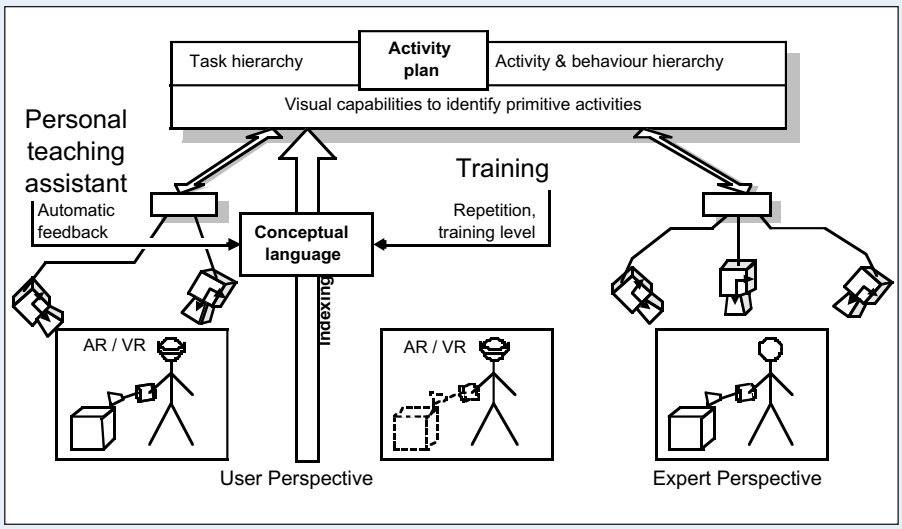

Using the ActIPret system to record and retrieve activities.

(1) replay of the sequence from arbitrary viewpoints, depending on the training level or (2) use of the ActIPret system in the form of a personal teaching assistant. In this case, for a selected task, the activities carried out by the user are automatically compared with the activities of the expert and improvements are suggested. This results in more effective training, compared to repetition without feedback.

ActIPret has two main technical objectives: the design and evaluation of a cognitive vision framework that extracts and interprets activities, and the development of purposive visual processing and interpretation techniques to provide the required perceptual capabilities.

To achieve a robust interpretation of activities, the interaction of visual attention, active camera behaviour, recognition, understanding, and knowledge from models, tasks, and context are being investigated. The interaction of these modules is the essential mechanism for removing possible ambiguities from the inherently uncertain information obtained through visual processing. The cognitive vision framework makes it possible to discriminate between activities that are essential to the task at hand (and should therefore be maintained) and those that are irrelevant (and should therefore be eliminated from the training sessions). The final outcome of the cognitive approach is the activity plan, which contains an index into activities and behaviours for access in user-driven training and for feedback while the trainee is rehearsing the activities.

To achieve the cognitive ability of the framework, vision techniques must provide the required functionality in the form of self-contained, cooperating components. The framework consists of both top-down (task-/behaviour-/context-driven) and bottom-up (data-driven, self-evaluating) interacting components. There are four types of visual processing components:

- extraction of cues and features
- detection of context-dependent relationships between cues/features
- recognition of activities and objects handled, taking into account potential occlusion
- synthesis of behaviours and tasks that modify the context of the other components.

All four types of components report visual evidence with confidence measures. These levels of visual interpretation are interlaced with the attentive and investigative behaviours that provide the feedback to purposively focus processing. Robust interpretation results are achieved with methods that actively seek out desirable viewpoints and obtain elucidative information for detection, recognition and synthesis. Robustness is also enhanced using context-dependent information integration between the components.

The ActIPret consortium consists of the following partners: Institute of Automation and Control, Vienna University of Technology (Project Coordinator), Center for Machine Perception at the Czech Technical University, School of Cognitive and Computing Sciences at the University of Sussex, UK, Computational Vision and Robotics Laboratory, ICS-FORTH and PROFACTOR - Produktionsforschungs GmbH, Austria.

Link:
Project web site:
http://actipret.infa.tuwien.ac.at/

Please contact:
Stelios Orphanoudakis,
Tel: +30 2810 391605,
E-mail: orphanou@ics.forth.gr

Antonis Argyros
Tel: +30 2810 391704,
E-mail: argyros@ics.forth.gr

Markus Vincze, Project Coordinator,
Tel: +43 1 5041446 11
E-mail: vincze@acin.tuwien.ac.at

Intelligent Traffic Light Control

by Marco Wiering

Growing numbers of road users and the limited resources provided by current infrastructures lead to ever increasing traveling times. The Intelligent Traffic Light Control project pursued at Utrecht University aims at diminishing waiting times before red traffic lights in a city.

Traffic in a city is very much affected by traffic light controllers. When waiting for a traffic light, the driver looses time and the car uses fuel. Hence, reducing waiting times before traffic lights can save our European society billions of Euros annually. To make traffic light controllers more intelligent, we exploit the emergence of novel technologies such as communication networks and sensor networks, as well as the use of more sophisticated algorithms for setting traffic lights. Intelligent traffic light control does not only mean that traffic lights are set in order to minimize waiting times of road users, but

also that road users receive information about how to drive through a city in order to minimize their waiting times. This means that we are coping with a complex multi-agent system, where communication and coordination play essential roles. Our research has led to a novel system in which traffic light controllers and the behaviour of car drivers are optimized using machine-learning methods.

Figure 1: Optimal control of traffic lights.

Our idea of setting a traffic light is as follows. Suppose there are a number of cars with their destination address standing before a crossing. All cars communicate to the traffic light their specific place in the queue and their destination address. Now the traffic light has to decide which option (ie, which lanes are to be put on green) is optimal to minimize the long-term average waiting time until all cars have arrived at their destination address. The learning traffic light controllers solve this problem by estimating how long it would take for a car to arrive at its destination address (for which the car may need to pass many different traffic lights) when currently the light would be put on green, and how long it would take if the light would be put on red. The difference between the waiting time for red and the waiting time for green is the gain for the car. Now the traffic light controllers set the lights in such a way to maximize the average gain of all cars standing before the crossing. To estimate the waiting times, we use 'reinforcement learning' which keeps track of the waiting times of individual cars and uses a smart way to compute the long term average waiting times using dynamic programming algorithms. One nice feature is that the system is very fair; it never lets one car wait for a very long time, since then its gain of setting its own light to green becomes very large, and the optimal decision of the traffic light will set his light to green. Furthermore, since we estimate waiting times before traffic lights until the destination of the road user has been reached, the road user can use this information to choose to which next traffic light to go, thereby improving its driving behaviour through a city. Note that we solve the traffic light control problem by using a distributed multi-agent system, where cooperation and coordination are done by communication, learning, and voting mechanisms. To allow for green waves during extremely busy situations, we combine our algorithm with a special bucket algorithm which propagates gains from one traffic light to the next one, inducing stronger voting on the next traffic controller option.

We have implemented the 'Green Light District', a traffic simulator in Java in which infrastructures can be edited easily by using the mouse, and different levels of road usage can be simulated. A large number of fixed and learning traffic light controllers have already been tested in the simulator and the resulting average waiting times of cars have been plotted and compared. The results indicate that the learning controllers can reduce average waiting times with at least 10% in semi-busy traffic situations, and even much more when high congestion of the traffic occurs.

We are currently studying the behaviour of the learning traffic light controllers on many different infrastructures in our simulator. We are also planning to cooperate with other institutes and companies in The Netherlands to apply our system to real world traffic situations. For this, modern technologies such as communicating networks can be brought to use on a very large scale, making the necessary communication between road users and traffic lights possible.

Please contact:
Marco Wiering
Utrecht University, The Netherlands
Tel: +31 30 253 9209
E-mail: marco@cs.uu.nl

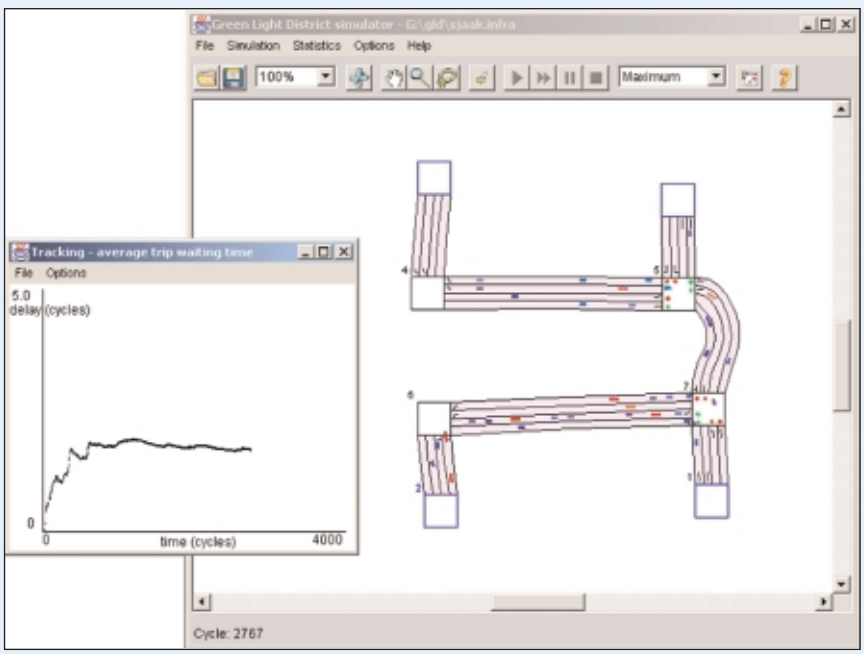

Figure 2: The simulator showing the infrastructure, road users, and plots of average waiting times.

Articles in this Section

"What did the others do?"

by Åsa Rudström and Petra Fagerberg

In real life, humans are social creatures, interested in the behaviour of others. We follow trails, we go where the action seems to be, and we buy the same kinds of things that our friends buy. The SOCIFER project at SICS investigated the effects of transferring some social mechanisms into the design of an on-line travel booking system. By getting statistical information on fellow travellers' bookings, users could see what others were doing and could use this as a basis for their own choices.

Recommendation is one way in which social functionality is exploited in interactive systems, eg on the Web. The recommendation may come from a specific person (such as a movie reviewer), be compiled from sales records (such as top-ten lists), or be based on the collective behaviour of many users (collaborative recommendation). An example of the latter is the 'people who bought this also bought that' functionality of Amazon.com.

The most prominent trend in recommender systems is to suggest what item to choose, or what action to take. However, for recommendations to really work it is crucial that the reasoning behind the recommendation be understood. When a friend recommends a book to you, you know that the recommendation is based on your friend's knowledge about you. A top-ten list is based on statistics, but if a booking system recommends that you add a cabin to your ferry ticket booking, why is this suggested to you?

In a joint project with Stena Line, an international transport and travel service company and one of the world's largest ferry operators, we explored the approach of visualising the collaborative data, leaving it to the user to decide on a course of action. In other words, instead of the letting the system make a suggestion, users were expected to create recommendations themselves based on the background data.

We developed and tested a prototypical social interface component, displaying fake but realistic statistical booking information concerning a limited number of ferry departures. A qualitative laboratory study was performed. Seven study subjects were given two ticket-booking tasks, one at a time, in varying order. The tasks were presented as mini-scenarios, giving most but not all of the information needed to accomplish the task. Subjects were minimally instructed and asked to 'think aloud'. Observations were made on user behaviour, and in particular when and to what extent the social component was noticed and used.

Study subject in action. Test leader to the right, technical supervisor in the background.

All subjects were overwhelmingly positive towards the idea of supporting the on-line booking process with relevant social information. We found strong evidence that subjects did use the statistics as a basis for their selections in the on-line booking. For example:

A: *"What about the other departure then? Not very many children, lots of noisy people, nah, we'll take the early boat."*
B: *"What did the others book here?"*
[talking about cabins]

We also found support for our hypothesis that users would form their own decisions based on background information. Statistics were used to decide both for and against booking choices:

C: *"… the others didn't book cabins, we don't need that either."*

D: *"… 'meal booking', no we won't do that. What did the others do? Brunch, aha. Maybe we should have that then …"*

The statistics also promoted some booking choices that were less visible in the interface:

E: *"Aha, breakfast… where do I book that?"*

None of the subjects in our study felt disturbed by the booking statistics being displayed, and they experienced very few problems in understanding the information.

We set out to test the hypothesis that an on-line booking system would benefit from making other people's bookings visible in the form of statistics. In conclusion, the results from the study support this hypothesis.

This work was performed during a period of nine months and ended in December 2002. The project was jointly funded by the Social Computing program of SITI (the Swedish Research Institute for Information Technology) and by Stena Line IT AB.

Link:
SOCIFER web site:
http://www.sics.se/humle/projects/

Please contact:
Åsa Rudström, SICS
Tel: + 46 70 774 8832
E-mail: asa@sics.se

Digital Image Processing in Art Conservation

by Barbara Zitová and Jan Flusser

Part of the 'Last Judgement' Mosaic conservation project has employed up-to-date image processing techniques to enhance an old photograph of the mosaic, and in particular to compare the current state of the mosaic with its condition in the nineteenth century. This work was done by the Institute of Information Theory and Automation at the Academy of Sciences of the Czech Republic, along with the Getty Institute, U.S.A. and the Prague Castle Administration.

In 1992-2000, the Getty Institute (Los Angeles, CA), jointly with specialists from the Czech Republic, realised The Last Judgement Mosaic conservation project. The mosaic, which is 84 square metres in area, was created in 1371 and is situated on the outer wall of the St.Vitus Cathedral in Prague, in the Czech Republic (see Figure 1). The dilapidation of the mosaic during the centuries was very severe, due to high levels of impurities in the mosaic glass, temperature fluctuations (the measured temperature range is -28°C to +60°C) and, more recently, air pollution. The conservation followed the original patterns and artist's intentions as closely as possible, and used all available historical sources. Unfortunately, it was only after the conservation had been finished that a photograph of the mosaic dating from the end of the nineteenth century was found. Nevertheless, this at least made a comparison possible between the restoration and the documented state of the mosaic at that time.

Figure 1: The St.Vitus cathedral in Prague Castle, Prague, Czech Republic. The Last Judgement mosaic is situated above the three arcades on the right.

The Department of Image Processing was invited to join the project and to use their expertise in advanced digital image restoration and registration (alignment) methods to reveal original patterns which had disappeared from the mosaic and which, while captured by the historical photograph, are obscured in it by the high level of noise and blurring. Moreover, they were asked to localise possible differences between the current and past states of the mosaic (1879).

For the experiments, a digital camera image of the current mosaic and the digitised negative (1200 dpi) of historical photograph were used. The quality of the old photograph is very low due to damage caused by aging, improper handling, etc. The special tonality of the historical image is caused by the low sensitivity of the original photographic material to the red part of the visible spectrum, so there were nontrivial differences between the intensity values in this photograph and in contemporary images. What made the evaluation of possible changes most difficult, however, was the difference in the geometry of the photographs, which had been taken from different angles and using different cameras. This fact forbade the use of direct image comparison for the change detection.

For the quality improvement of the historical photograph, several image restoration techniques were applied. We used various noise suppression methods (wavelet-based methods, adaptive non-

Figure 2: Examples of differences identified in the part of the mosaic called 'Resurrection' between the current (left) and nineteenth century versions (right): (1) there is an extra wave in the headdress (2) different shape of the ornamental pattern (3) the horizontal coffin edge is missing.

linear filters, Mumford-Shah's denoising), image sharpening (total variation sharpening, inverse heat equation approach, Wiener filter) and combined restoration techniques (total variation reconstruction with Gaussian point spread function estimation). The latter methods gave the best results.

The photographs were brought into geometrical alignment by means of an appropriate registration technique. Besides the classical feature-based method, in which the correspondence of the detected feature points is used for computation of the transform parameters, the mutual information similarity measure was used for improvement of the feature points localisation. This method is designed specifically for the case of data acquired by different sensors. Thus, the non-correspondence of intensities in the photographs did not affect the quality of the final alignment of the photographs. The differences between photographs were then localised. Several of them were apparent by simple visual inspection, but others appeared only after the exact alignment and analysis of the overlaid transformed photographs. This would be impossible to achieve without the help of digital image processing (see Figure 2).

The presented part of the project illustrates that digital image processing can be a useful tool, applicable in the area of art protection and conservation, and is able to supply information not otherwise available. Nevertheless, because of the complex and highly interdisciplinary nature of the problems arising in this area, it is unlikely that digital image processing techniques will become completely automated. Human assistance - computer scientists as well as art historians - will always be required.

Link:
http://www.utia.cas.cz/user_data/zitova/research.html

Please contact:
Barbara Zitova, Jan Flusser
Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic/ CRCIM
Tel: +420 2 6605 2357
E-mail: {zitova, flusser}@utia.cas.cz

A Real-Time Infrared Tracking System for Virtual Environments

by Maxim Foursa

Tracking, also called position and orientation tracking or position tracking and mapping, is used in Virtual Environments where the orientation and the position of a real physical object is required, for example to interact directly with the 3-D objects rendered by the virtual environment system. An optical tracking system has been developed in the laboratory of Virtual Environments (VE) of the Fraunhofer Institute for Media Communication in collaboration with the Russian Institute of Computing for Physics.

The optical tracking system operates on a standard PC platform and includes three cameras and a frame grabber. The system is able to track the position and orientation of a user's head and stylus in a CAVE-like volume. Active LEDs generate optical markers, which are acquired by the cameras. The system follows events in real time, processing 25 image frames per second. The estimated latency is below 100ms and the measured accuracy of reconstruction of marker position in space is below 5mm. Since this system uses commodity hardware and internally developed software, the price of the system is around 4000 Euro for the the configuration with three cameras and 3300 Euro for the configuration with two cameras.

Specifying a point in three dimensions requires the transition position, that is, the Cartesian coordinates x, y, and z. However, many VE applications manipulate entire objects and this requires the orientation to be specified by three angles known as pitch (elevation), roll, and yaw (azimuth). Thus, six degrees of freedom form the minimum required to fully describe the position of an object in 3-D.

Trackers are used to measure the motion of the user's head or hands, and sometimes eyes. This information is then used to calculate the correct perspective projection and auditory input according to the user's position.

Tracking may be performed using a number of different methods: there exist magnetic, mechanical, optical, acoustic (ultra-sonic), and inertial head trackers. These types of trackers can also be

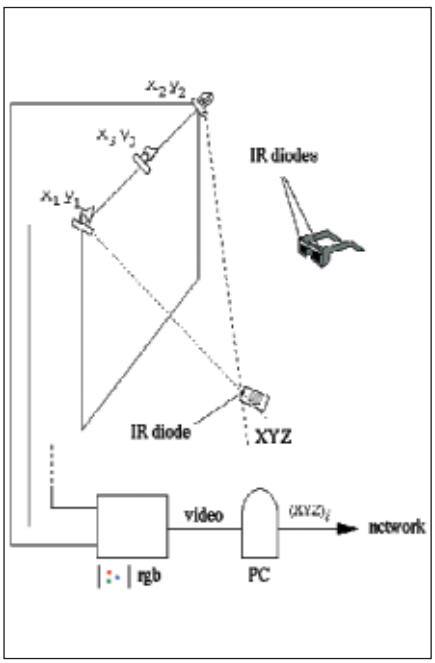

Figure 1: System setup.

mounted on glove or body-suit devices to provide tracking of a user's hand or some other body part.

For example, in the case of magnetic sensors, a receiver is placed on the user's head so that when the head moves, so does the position of the receiver. The receiver senses signals from the transmitter, which generates a low-frequency magnetic field. The user's head motion is sampled by an electronic unit that uses an algorithm to determine the position and orientation of the receiver in relation to the transmitter.

With almost all of these systems, the user has to be linked to a measurement instrument, via either a cable or, even more restrictive for the user, a mechanical linkage. Furthermore, while mechanical tracking systems are extremely precise, magnetic and acoustic tracking systems suffer from distortion originating from a number of sources. Furthermore, since magnetic systems use magnetic fields, which are generated proximal to the user, it is not safe for people to spend extended times using them. For this reason, we started developing an optical tracking system, which should overcome many of the drawbacks of conventional tracking systems.

Optical tracking systems can work fast over a large area and be comfortable to use, but are limited by light source intensity and require a line of sight from emitter to receiver. Though high-quality, precise optical systems are quite expensive, we have demonstrated that it is possible to make a simple system that meets the requirements of user-VE-user interaction, on the basis of commodity hardware and complex multi-step software processing..

Our system employs three near-infrared monochrome cameras JAI-M50IR, which are equipped with 6mm lenses and infrared filters and are installed in a test laboratory above a reflective screen. These cameras are attached to an RGB frame grabber. The output from the frame grabber is one image with the size 769x570 that contains synchronised output from all the cameras in the R, G and B planes. The scanning speed of the frame grabber is 25 frames/sec. This

Figure 2: Devices to track.

Figure 3: User with the tracking devices.

output is processed on a workstation with Pentium 4 2.2GHz processor, using specially developed software. To resolve the objects being tracked, we are using active infrared beacons. Currently we have three devices to track: a pointer with one LED emitter, shutter glasses with two LEDs and a pointer with three LED emitters and a telescopic tube, for better direction reconstruction. The power supply for the devices is a 9V battery. The direct current is constant and can be varied from 5mA to 50 mA.

The general scheme of operations is the following. The first step after the frame grabber acquires the three images is the beacon detection. It can operate in two different modes. The first mode is a global search over the whole image data in one colour plane (that is the output from one camera) for a detection of all beacon positions. If the system knows the last two calculated positions, the second mode will do a local search because a prediction has restricted the area of interest. When we have all beacon positions on the image from one camera, we can calculate the epipolar line and optimise the search process on other cameras by doing it in the vicinity of the line. To get the coordinates of a beacon, the program looks for a light spot ellipse (the motion causes an ellipse rather than a circle) and determines its radii. The next step is a 2-D transformation from distorted image coordinates to undistorted camera sensor coordinates. Afterwards, the epipolar constraint can be used to get the correlated image points, as well as other constraints defined by beacon positions or sizes, and as a result, we obtain the 3-D positions which we can assign to corresponding devices.

For calibration we use a two-step method: to get the internal camera parameters we use an implemetation of the Tsai method, and to get its position and orientation we use the Faugeras method and a special target which has more than thirty points measured by a theodolite with high accuracy in a coordinate system.

The system is currently installed in the CyberStage VE system at Sankt-Augustin and is being used instead of or with the Polhemus FASTRACK magnetic tracking system. The CyberStage is a 3m x 3m CAVE with stereoscopic projections on all four walls and the floor, and with 8-track sound and vibrating floor. The tracking server transmits the data via the network to a listening daemon, which processes them for AVANGO programming framework, which then renders the virtual scene. The system is flexible and can be adjusted for the use in other virtual environments, if the lighting conditions (eg the absence of bright light, such as sunlight or the light emitted by CRT monitors) allow it. It is anticipated that the system will be used in the Russian Institute of Computing for Physics and Technology in low-cost VR projects.

Please contact:
Maxim Foursa, FhG-IMK
Tel: +49 2241 12 2773
E-mail: maxim.foursa@imk.fhg.de

Content-based Retrieval Services in Peer-to-Peer Systems using Taxonomies

by Yannis Tzitzikas and Carlo Meghini

We investigate an approach where the participants of a peer-to-peer system use taxonomies to describe the contents of their objects and to formulate queries to each other. Inter-taxonomy mappings are employed in order to carry out the required translation tasks.

There is a growing research interest in peer-to-peer systems, like Napster, Gnutella, FreeNet and many others. A peer-to-peer (P2P) system is a distributed system in which participants rely on each other for certain services, rather than relying solely on dedicated and often centralized infrastructures. The membership in a P2P system is relatively unpredictable because it is ad-hoc and dynamic: services are provided by the peers that happen to be participating at any given time. Several examples of P2P systems have emerged recently, most of which are wide-area, large-scale systems that provide content sharing, storage services, or distributed 'grid' computation. Mainly they focus on specific applications (eg music file sharing) or on providing file-system-like capabilities. They do not yet provide content-based retrieval services: in most of the cases, the name of the object (eg the title of a music file) is the only way to describe the contents of the objects.

In general, the language that could be used to describe content and to formulate content-based queries is either free (eg natural language), or controlled, ie object descriptions and queries may have to conform to a specific vocabulary and syntax. The former case resembles distributed Information Retrieval (IR) systems and this approach is applicable when the objects of an application have a textual content. Our research focuses on the latter case where the objects of a peer are indexed according to a specific conceptual model which may be represented by various data models (eg relational, logic-based, etc), and content searches are formulated using a specific query language. Of course, a P2P system could impose a single conceptual model on all participants to enforce uniform, global access, but this would be too restrictive. Alternatively, a limited number of conceptual models could be allowed, so that traditional information mediation and integration techniques will likely apply (with the restriction that there is no central authority). The case of fully heterogeneous conceptual models makes uniform global access extremely challenging.

The first question to be investigated is which conceptual modelling approach is appropriate for P2P systems. We need an approach which makes it possible to bridge various kinds of heterogeneity in a systematic and easy manner. As there are no central servers, or mediators, the participating sources must have (or be able to create) articulations, ie, mappings, between their conceptual models in order to be able to translate exchanged queries appropriately. Such mappings could be constructed manually but the more appropriate - and more challenging - approach for a P2P network is automatic articulation. For all these reasons, a simple, clear, and application-independent conceptual modelling approach appears advantageous.

Our research investigates an approach that is based on taxonomies. Taxonomies

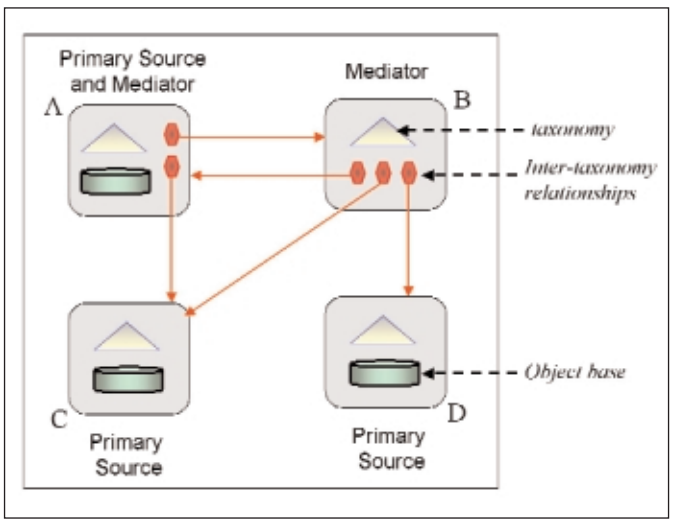

Figure 1: The architecture of a P2P system using taxonomies and inter-taxonomy relationships.

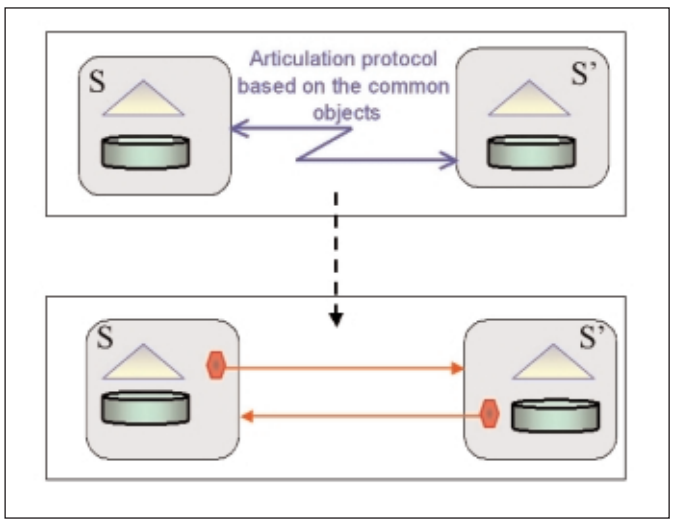

Figure 2: Data-driven automatic articulation.

are relatively easy to build in comparison with other kinds of conceptual models. They can be constructed from scratch or can be extracted from existing taxonomies (eg from the taxonomy of Yahoo!) using special-purpose languages and tools. Data-driven methods for taxonomy mapping can be automated and are therefore more appropriate for P2P networks. These methods can be used to create mappings between two taxonomies on the basis of the objects that are indexed by both taxonomies. According to our setting, a source in a P2P system can serve any or all of the following roles: primary source, mediator, and query initiator. As a primary source it provides original content to the system and is the authoritative source of that data. As a mediator it does not store any content: its role is to provide a uniform query interface to other sources. As a query initiator it acts as client in the system and poses new queries. Figure 1 shows the architecture of a network consisting of four peers A, B, C and D; two primary sources (C and D), one mediator (B) and one source that is both primary and mediator (A).

Apart from the classical problems that are currently being studied in the area of P2P networks with respect to query evaluation, such as object placement, replication, caching and freshness, the support of content-based retrieval services raises new questions. For instance, query forwarding requires techniques for query translation. Although this can be done using the techniques employed by mediators, the fact that in P2P systems we can have mutually articulated mediators means that endless-query-loops may arise and cause dead-lock phenomena. We are now studying techniques in order to avoid such phenomena and to optimize query evaluation.

The first author currently holds an ERCIM fellowship. This work is result of research conducted at ISTI-CNR, Pisa.

Please contact:
Yannis Tzitzikas - ISTI-CNR
Tel: +39 050 3152901
E-mail: tzitzik@iei.pi.cnr.it
http://www.csi.forth.gr/~tzitzik

Geometric Numerical Methods for Continuum Dynamics

by Jason Frank

Qualitative properties of numerical simulations of continuum dynamics are studied at CWI using geometric integration. In particular atmospheric fluids were studied with particle-based numerical flow models, in co-operation with Imperial College of London. The ongoing research includes the simulation of ferromagnetic materials.

The project started in 2002, and is funded by the Netherlands Organization for Scientific Research NWO (Innovative Research Grant Veni). The work extends developments in the previous decade in the field of geometric integration, primarily focussed on ordinary differential equations, to continuum systems (fluids, solids) described by partial differential equations.

In many areas of scientific interest, numerical simulations are pushed beyond the limits of classical considerations in numerical analysis such as global error and stability. In molecular dynamics and celestial mechanics, for example, systems that are known to be sensitive to initial conditions are integrated numerically over very long time intervals, despite the fact that the global error has long reached 100 percent. In this case, the user is interested in exploring the qualitative features of the solution rather than in exactly reproducing the trajectory starting from a given initial value.

In the late 1980s and throughout the 1990s, the field of geometric integration arose to develop and analyze numerical methods for ordinary differential equations (ODEs) with the above purpose. Figure 1 illustrates the difference between a geometric integrator and an ordinary method for an ODE of interest in atmospheric dynamics. The most important discovery in geometric integration is that, roughly speaking, a numerical method which retains some property, such as energy conservation or reversibility, actually produces the solution of a perturbed problem with that property. By analyzing the perturbed problem, one can make statements about the quality of the solution far beyond the range of global error analysis. (Hairer, Lubich & Wanner, Geometric Integration, Springer-Verlag, 2002).

Scientists studying continuum processes may also be more interested in the qualitative behaviour of a solution than in the exact solution starting from a particular initial state. One example is climate simulations, which may be carried out over intervals of 100 years or more. Currently, interest in geometric integration is shifting to partial differential equations (PDEs). In some sense, methods for PDEs have had 'geometric' qualities for a long time; the use of monotone methods and Riemann solvers for solving hyperbolic problems is a good example. However, there are new developments as well. CWI, in cooperation with Imperial College of London, has made considerable progress in the last two years in the development

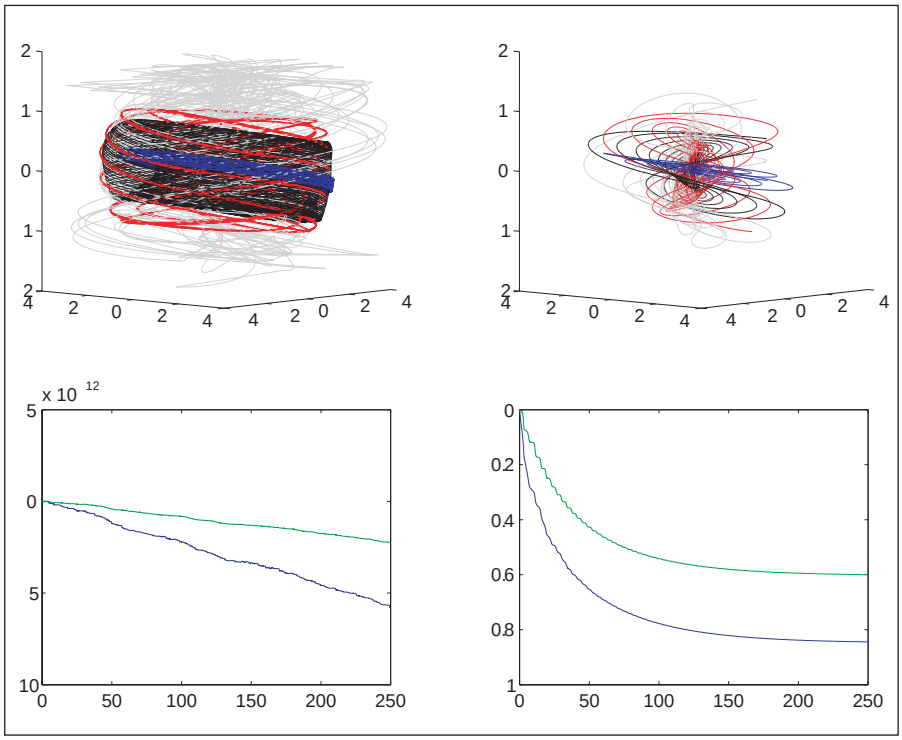

Figure 1: Solutions of the 5-component Lorenz model for 4 different initial states, distinguished by colour (top two plots), including one chaotic trajectory (in grey). This model studies geostrophic balance in atmospheric/oceanic fluids. The bottom two plots indicate change in total energy (blue) and enstrophy (green), both of which should be conserved in time. The left two plots are for a geometric integrator, which conserves energy and enstrophy to within 0.01 billionth. The right two plots are for an ordinary method, which damps out more than 80 percent of the energy, and introduces an artificial steady state.

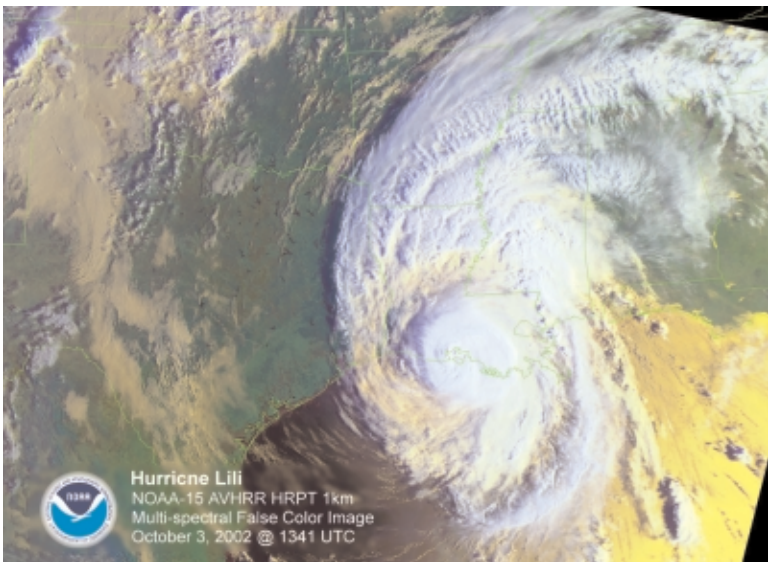

Figure 2: Vortical structures occur on all scales in the atmosphere, hurricanes being a familiar and spectacular example. The study of such structures is very important for our knowledge about the atmosphere. The numerical method of geometric integration, which conserves properties such as circulation, is particularly suited for this purpose (courtesy NOAA/National Climate Data Center, U.S.A.).

of geometric methods for atmospheric fluids using particle-based numerical fluid models. The Hamiltonian Particle-Mesh method combines finite mass fluid particles for advection with a grid for fast evaluation of derivatives, such that the resulting method conserves energy and circulation (vorticity) in addition to mass.

Other research at CWI on geometric integration concerns the simulation of ferromagnetic materials. In particular, the equations governing the flow of magnetic orientation in such materials implicitly define a set of conservation laws for energy and momentum. A numerical method applied to the equations of motion would normally not retain any discrete analogue of such conservation laws. We are currently investigating a method with precisely this property.

Please contact:
Jason Frank, CWI
Tel: +31 20 592 4096
E-mail: jason@cwi.nl
http://www.cwi.nl/~jason/

Figure 1: Surface to reconstruct.

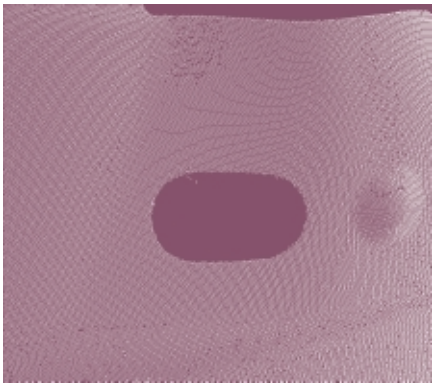

Figure 2: Input point-set.

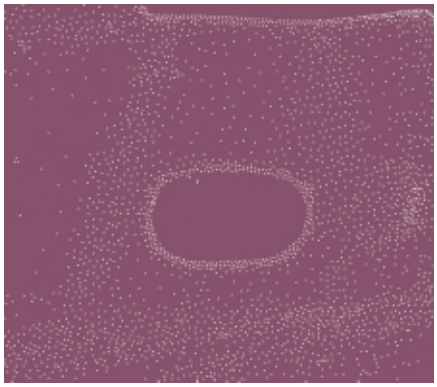

Figure 3: Filtered points.

Surface Reconstruction - Triangulating Scattered Points

by Géza Kós

Triangulating surfaces from unorganised sets of sampling points is an important problem in the reverse engineering of geometric shapes. At the Geometric Modelling Laboratory of SZTAKI, a new, efficient triangulation algorithm has been developed.

Reverse engineering of geometric shapes is the process of converting a large number of measured data points into a concise and consistent computer representation. In this sense, it is the inverse of the traditional CAD/CAM procedures, which create physical objects from CAD models.

Triangulating scattered point-sets is a very important problem of reverse engineering. Given a set of unorganised points that lie approximately on the boundary surface of a three-dimensional object, and without a priori information on the topology, our goal is to reconstruct the surface by building a triangular mesh using the given points as vertices.

The resulting polyhedron can be the input of further procedures like surface fitting, or can be visualised with various textures. (For example, in computer-animated films the characters are often created as clay models first, then the 3D scanned and triangulated models are used for visualisation.)

All algorithms aiming to solve this problem must overcome several difficulties. The first one is related to the size and quality of the input. Modern 3D scanners make it possible to acquire several (ten) millions of sample points on the object's surface. For current applications, a triangulation algorithm must be prepared to efficiently handle such huge data sets. Furthermore, the point density is often very uneven due to curvature variations, undesirable outlying elements may occur, and scanning techniques affect sampling distribution. The input also contains noise (measurement errors).

Another problem is related to the topology of the surface. The measured object can be of arbitrary topology, and its surface may have positive Genus. It is also typical that the point-set represents not a complete volumetric object, but only certain surface portions of the boundary, and only these parts need to be reconstructed. In such cases, the surface has holes and/or open boundaries.

The main algorithmic difficulty is related to the dimension. Due to Voronoi's and Delaunay's works, creating tessellation in the space of any dimension is very easy if the tiles have the same dimension as the space. The triangulation problem is slightly different, since the algorithm must create a two-dimensional object in a three-dimensional space. This leads to more complicated topological situations.

Previous algorithms have used many interesting ideas. Some basic algorithms work by projecting the points to a carrier surface and creating triangulation in the parametric domain of the carrier surface. These methods are somewhat limited for disconnected surface portions and objects with positive Genus. Other approaches start from various subsets of the 3D Delaunay tessellation of the sampling points and try to select a proper subset of the faces. Another very popular idea is to define implicit surfaces containing the given sampling points and extract these surfaces.

The new algorithm developed in SZTAKI offers a different approach to solving the problem. Our algorithm considers the surface as a metric, 2D manifold. Using the position of the sampling points, distances and angles within the manifold are estimated and the relative Delaunay triangulation of the sampling points is computed. The triangulation is built locally around each data point and these local triangulations are merged to give a consistent mesh representation of the surface.

The algorithm was found to be efficient. The current implementation is capable of processing huge data sets: up to 100,000,000 vertices on a commercial

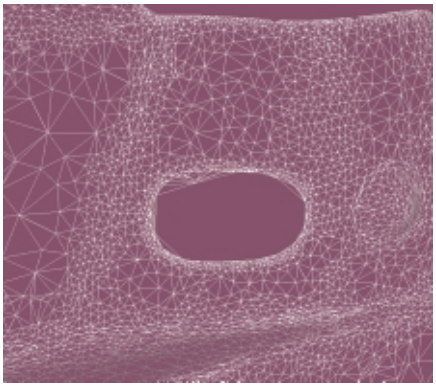

Figure 4:
Triangulation on the filtered points.

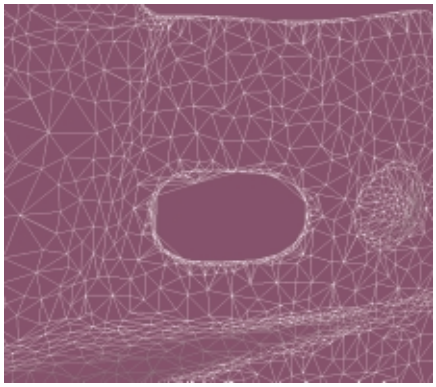

Figure 5:
Decimated triangulation.

desktop PC, which generates up to 200,000,000 triangles.

Since most applications require smaller data sets, the algorithm has been enhanced with two optional decimation steps which reduce the size and complexity of the data. The first decimation step is the curvature-based prefiltering of points, which can be executed before generating any triangle. This procedure eliminates many points in the large, smooth areas of the surface and retains the original density in the neighbourhood of highly curved regions and boundary loops. By this step, the number of points can be reduced by 70-90 percent, without significant loss of quality in the final result.

The second decimation step is performed after generating the triangles. This procedure takes advantage of the fact that triangulation makes a better error estimate possible and more vertices can be deleted.

This project started within the framework of an EU-supported COPERNICUS project (RECCAD no. 1068) in 1997, and has also been supported by the Hungarian Ministry of Education, Grant No. OMFB-01989/2002.

Link:
SZTAKI GML web page:
http://www.sztaki.hu/gml

Please contact:
Géza Kós, SZTAKI
Tel:+36 1 386 8782
E-mail: kosgeza@sztaki.hu

PROXiTV: A Multimedia Distribution Infrastructure

by Éric Fleury and Isabelle Guérin Lassous

Access to audio/video sequences via Internet is a frustrating experience both in terms of interactivity and of quality of images and sound. The main objective of the PROXiTV project was to develop a proxy infrastructure to deliver high quality multimedia content to consumers connected to broadband local loops (ADSL).

PROXiTV aims to establish a high bandwidth Internet solution and interactive TV infrastructure for high quality TV delivery over IP networks, by taking advantages of the existence of local broadband loops. The basic idea is to build a global technical solution for providing really interactive services to a large public of home users connected directly, from their PC or TV, to a proxy server at the head traffic point of broadband local loops. In order to reach these ambitious objectives, PROXiTV addressed several issues: (1) Build an efficient distribution infrastructure of multimedia proxies; (2) Develop new interactive content-oriented services exploiting true bandwidth access;

The PROXiTV consortium gathers two INRIA research teams (REMAP and RESEDAS); Jet2Web/Telekom Austria which is the dominant telecommunications operator for fixed network services and the leading Internet Service Provider in Austria; Webfreetv.com a full-service interactive TV broadcaster with special field of corporate communication; Eurosport, a pan-European thematic sports channel broadcast via cable and satellite systems and EADS-Sycomore, a French company involved in the development of Information Systems including video solutions, real time TV systems and high performance servers design.

The experimentations done in PROXiTV represent the most important part of the project. Indeed, the main issue was to conduct a large scale experiment on 3 ADSL local loops (Vienna, Linz and Steyr), including B-to-C services with high quality interactive Web (with TV and VOD), and B-to-B applications. The services tested during the experimental period were: Music clip (in cooperation with EMI Records, Virgin Records and Universal Music); Digital Video Recorder (customers have the possibility to record content form the Live TV part on the Telekom service); Videoconferencing application. The use of the services (the monthly basic subscription fee and all pay by view) is billed directly over the Telekom Austria bill. The expected results were to give an estimate if interactive TV over IP-Networks makes sense in terms of

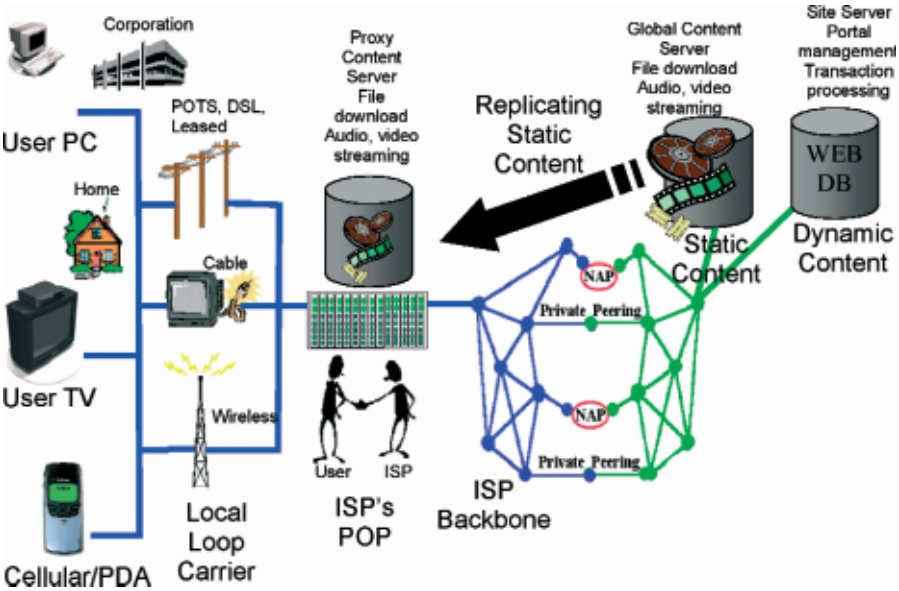

Figure 1: Replicating static content closer to the user.

network resources and from an economical perspective.

The PROXiTV infrastructure developed in the project is a global solution to deliver videos to large groups of users with a high quality of services. It is made of different parts. The server part is based on the PeakServer solution provided by EADS-Sycomore. PeakServer is a server platform designed for on-demand distribution of Multimedia contents on broadband networks from simple web pages delivery to broadcast quality video streaming. PeakServer can deliver media-rich web pages, stream high quality video contents (NVOD, VOD) or download various contents such as games, software, music or movies. The CDN (Content Delivery Network) is based on the Peakframe solution developed by EADS-Sycomore for the Telecom and Media business with the purpose to commit on an end-to end Content Delivery technical Infrastructure addressing commercial services on broadband portals. The global PeakFrame/Peakserver solution includes: · PeakFrame EntryPoint for contents acquisition, content management, audience report and supervision. The route and transfer optimization of the content traffic needed to update all the ProxiTV Content Servers from the PeakFrame Entry Point was conducted by INRIA. The solution consists in an application level multicast tree in order to replicate data from the PeakFrame Entry Point to all Proxy Content Servers. The application level multicast tries to optimize several criteria as bandwidth, delay, etc.

Figure 2: Portal with remote control and set top box.

Now at the end of PROXiTV project, which has given us the chance to look over the garden fence into the world of interactive TV (iTV), it is clear that we are at the early beginning of a totally new service. Unfortunately we did not have yet information about set-top box usage, which might be five or ten times higher as the usage of an PC user. Nevertheless it seems quite clear that multimedia and iTV will be an appropriate solution to solve problems in the Internet access area. Internet access in the narrowband field is saturated. PROXiTV project is one of the leading iTV projects and the most important result of the Project is "there is a real demand on such services".

Link:
http://www.istproject.com/proxitv/

Please contact:
Eric Fleury, INSA-Lyon Scientific and Technical University and INRIA
Tel: + 33 472 436 421
E-mail: Eric.Fleury@inria.fr

VTT inaugurated Microelectronics and Nanotechnology Centre 'Micronova'

by Pia-Maria Linden-Linna

Micronova is a joint project of the VTT and the Helsinki University of Technology (HUT), The nationally unique and internationally competitive research cluster will focus on microelectronics and nanotechnology research and innovative activities based thereon. The specially designed clean room facilities that are required in research are the largest and most highly rated in Scandinavia, and among the five largest in Europe.

MICRONOVA, was officially inaugurated on 4 March 2003 in Espoo Finland. Distinguished speakers at the ceremony were Erkki K.M. Leppävuori, Director General of VTT, Paavo Uronen, Rector HUT, David L. Tennenhouse, Vice President Intel Corporation, Juhani Kalanti, Director STMicroelectronics and Asko Vehanen, Executive Vice President, Okmetic Oyj.

Micronova is by far the largest research centre in the area of microelectronics, solid logic technology and sensor technology in Finland, which even at present has a clientele covering the whole country. The centre enables improved connections between basic research and applied industrial research and product development. VTT also has a highly rated research unit with clean room facilities in Oulu, the north of Finland.

Micronova clean room facilities.

The rapid development of microelectronics has for a long time maintained the growth rate of the electronics industry at double that of industry as a whole. The aim for smaller, faster, cheaper and less energy consuming equipment increases the significance of research in the area.

"The development of solid logic technology and materials will continue into the foreseeable future. This development is the foundation, on which the force that changes the IT world is built. The development trend is stable; fluctuations in economic market conditions only have a marginal effect after all. Microelectronics is becoming an integration platform for various technologies. It is no longer purely electronic development; instead, the functions of microsystems built on solid logic chips are activated by not only electronic and magnetic impulses, but also mechanical, chemical and biological impulses. Several companies find business opportunities from this technological platform", emphasised VTT Information Technology's Executive Director Pekka Silvennoinen.

"Information technology and electronics will continue to be a growing branch of industry, even though the economic outlook has decreased considerably since the construction decision was made. The demand for electronics research grew throughout the 1990s despite the recession", said VTT's Manager of Enabling technologies Ilkka Suni.

Micronova also has potential on the global market, where IT is growing more rapidly than in Finland. The units located in the centre have a broad international co-operation network.

Micronova emphasises the utilisation of new technology. The centre functions as a microelectronics business park for companies that require a high quality product development environment, offering the use of Micronova's facilities, research equipment and expertise. The model of operations offers companies a cost-efficient opportunity to utilise large public investments.

Three research areas of VTT Information Technology, units of various HUT departments and five client companies are represented in Micronova. Micronova provides facilities for a total of approximately 320 people.

Including production equipment, the Micronova complex represents an investment of almost • 100 million. The construction investment for the now completed expansion phase was • 33.6 million, and the total investment for both phases (1997/2002) was • 50 million. The net area of the new expansion is 10,689 square metres. Micronova's total net area is approximately 15,000 square metres, of which 2,600 square metres is high quality clean room.

Please contact:
Pia-Maria Lindén-Linna
VTT Information Technology
Tel: +358 9 456 4501
E-mail: pia-maria.linden-linna@vtt.fi

euroCRIS
Current Research Information Systems

Seminar on Current Research Information Systems

Brussels, 19-20 May 2003

A seminar on CRIS (Current Research Information Systems) will take place on 19-20 May 2003 in Brussels. Organised by *euroCRIS*, it has attracted the active partner participation from the following pan-European organisations:

- ALLEA - All European Academies
- CODATA - Committee on Data for Science and Technology
- EARMA - European Association of Research Managers and Administrators
- ERCIM
- ESF - European Science Foundation

A preparatory meeting with most of the partners in this enterprise was held in Antwerp 19 February (see http://www.ub.uib.no/avdeling/fdok/cris/antwerp03/partnermeeting.htm

euroCRIS is a not-for-profit association, established in Europe to be the internationally recognized point of reference for all matters relating to CRIS. A CRIS is any information tool dedicated to provide access to and disseminate research information. This includes people, projects, organizations, results (publications, patents and products), facilities, equipment, etc. To see examples, consult the NIWI-maintained DRIS world-wide directory of CRIS applications at http://www.niwi.knaw.nl/ cgi-bin/nph-dris_search.pl?language=us.

The primary goal of *euroCRIS* is to act as a single forum for all interested individuals and organizations to enter into dialogue and resolution of all matters related to the use of information technology in the conduct of all research information system business. *euroCRIS* (through its members) is responsible for maintenance of standards and provision of advice in all matters relating to CRIS. Additionally, *euroCRIS* supports standardized, streamlined information exchange across all aspects of the CRIS lifecycle as follows:

- apply and maintain internationally accepted standards and guidelines, and develop where necessary (eg Common European Research Information Format - CERIF)
- build and maintain an RTD thesaurus
- establish a 'one-stop shop' portal/ gateway to global CRIS
- set up a European database of RTD endeavors
- adhere to principles of best practice
- include the management of the Code of Good Practice (CGP)
- provide a forum for exploring and exploiting new and emerging concepts and technologies (incl. data quality, standards, etc…)
- nurture a community of CRIS practitioners and users
- addressing the needs of the CRIS user groups.

The Common European Research Information Format (CERIF) has developed since the late eighties as a recommendation to member states of the European Union. It is intended for use in CRISs (Current Research Information Systems). The European Commission has transferred the custodianship of CERIF to euroCRIS. Its purpose is to :

1. provide a design template and tools for a CRIS developer
2. enable exchange of R&D information between CRISs
3. enable portal access to R&D information in CRISs.

CERIF-compliant CRISs are of benefit to:

- the policymaker
- the R&D administrator in a funding agency
- the R&D manager in a university or government laboratory
- an intermediary such as a management consultancy
- a SME looking for patentable and wealth creating opportunities.

Successful CERIF-compatible CRISs exist currently in 7 countries and many more are under development.

More information:
http://www.eurocris.org/cerif

SPONSORED BY ERCIM

Twelfth International World Wide Web Conference

Budapest, 20-24 May 2003

The conference will start with a day of tutorials and workshops and followed by a three-day technical programme. The fifth day will be a 'Developers Day.' The tutorials and workshops will provide in-depth looks at specific areas of current interest. The technical programme will include refereed paper presentations, alternate track presentations, plenary sessions, panels and poster sessions. Developers Day will be devoted to in-depth technical sessions designed specifically for Web developers.

Further information:
http://www2003.org/

CALL FOR PARTICIPATION

7th International Workshop on Advanced Infrared Technology and Applications

Pisa, Italy, 9-11 September 2003

The Workshop, which is the seventh event of a series started in 1991, will constitute a forum bringing together scientists and young researchers, from both the academic and industrial worlds, to exchange knowledge, ideas and experiences in the field of infrared science and technology.

AITA 2003 will be held at the prestigious Scuola Normale Superiore of Pisa, sited in the Piazza dei Cavalieri. the central square of medieval Pisa, at a hundred meters from the marvellous Piazza dei Miracoli and the famous Leaning Tower.

The Workshop is organized jointly by the Fondazione Giorgio Ronchi and the Italian National Research Council.

More information:
http://ronchi.iei.pi.cnr.it/AITA2003/

ERCIM NEWS

ERCIM News is the magazine of ERCIM. Published quarterly, the newsletter reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of over 9000 copies.

ERCIM News online edition is available at http://www.ercim.org/publication/ERCIM_News/

ERCIM News is published by ERCIM EEIG, BP 93, F-06902 Sophia-Antipolis Cedex
Tel: +33 4 9238 5010, E-mail: office@ercim.org
ISSN 0926-4981

Director: Bernard Larrouturou
Central Editor:
Peter Kunz
peter.kunz@ercim.org
Local Editors:
AARIT: Wolfgang Hawlik
w.hawlik@ocg.at
CCLRC: Martin Prime
M.J.Prime@rl.ac.uk
CRCIM: Michal Haindl
haindl@utia.cas.cz
CWI: Henk Nieland
Henk.Nieland@cwi.nl
CNR: Carol Peters
carol@iei.pi.cnr.it
FORTH: Yannis Pachoulakis
yannisp@ics.forth.gr
FhG: Michael Krapp
michael.krapp@scai.fraunhofer.de
FNR: Patrik Hitzelberger
hitzelbe@crpgl.lu
INRIA: Bernard Hidoine
bernard.hidoine@inria.fr
NTNU: Truls Gjestland
truls.gjestland@ime.ntnu.no
SARIT: Harry Rudin
hrudin@smile.ch
SICS: Kersti Hedman
kersti@sics.se
SRCIM: Gabriela Andrejkova
andrejk@kosice.upjs.sk
SZTAKI: Erzsébet Csuhaj-Varjú
csuhaj@sztaki.hu
TCD: Ann McNamara
Ann.McNamara@cs.tcd.ie
VTT: Pia-Maria Linden-Linna
pia-maria.linden-linna@vtt.fi

Free subscription

You can subscribe to ERCIM News free of charge by:

- sending e-mail to your local editor
- contacting the ERCIM office (see address above)
- filling out the form at the ERCIM website at http://www.ercim.org/

SZTAKI - **First prize of the 'Agentcities Technology Competition and Exhibition'**. SZTAKI's System Development Department won the first prize of the 'Agentcities Technology Competition and Exhibition' in the Infrastructure Category for their WSDLTool, a tool for deploying web service wrapper agents. Agencities is an IST-funded initiative to construct an open network of on-line systems hosting diverse agent based services. 54 entries from all over the world had registered for the competition. The prizes were presented during an event in the Barcelona city hall on 6 February, featured on the Spanish national television. Fore more information, see http://www.agentcities.org/EUNET/ID3/

CCLRC - **Computational Science and Engineering Department at CCLRC, along with partnerspartners EPCC and IBM, are providing the flagship high performance computing service for UK academic research - known as HPCx.** The HPCx service is based on

IBM POWER 4 technology and provides an initial capability of 6.7 Teraflops/s peak. This performance will be upgraded to more than 11 Teraflops/s in 2004 and to 22 Teraflops/s in 2006. The HPCx system consists of 40 IBM pSeries 690 Regatta nodes, each containing 32 POWER4 processors, a total of 1280 processors. It is located at the UK's CCLRC Daresbury Laboratory and operated by the HPCx Consortium At the start of service on 9 December 2002 it became the most powerful high performance computer for academic research in use in Europe.

CWI - **Paul Vitányi** has been appointed CWI Fellow from January 1st, 2003. This

position is given to outstanding members of CWI's research staff, enabling them to concentrate completely on research. Vitányi joined CWI in 1971. Under his direction his research group, now led by Harry Buhrman, became a world leader in quantum computing, distributed algorithmics, algorithmic information theory, learning and inference methods, and reversible adiabatic computing. Together with Ming Li he pioneered applications of Kolmogorov complexity and co-authored 'An Introduction to Kolmogorov Complexity and its Applications', Springer-Verlag, New York, 1993 (2nd Edition 1997), parts of which have been translated into Chinese, Russian and Japanese. The Chinese version won the 1999 first prize for Excellent Books in Science and Technology published in China. Vitányi serves on the editorial boards of several internationally renowned journals. He is also a professor of Computer Science at the University of Amsterdam. Web page: http://www.cwi.nl/~paulv/

INRIA - **70 post-doctoral positions available in 2003**. INRIA will be recruiting approximately 70 post-doctorates, 23 of whom will be part of the annual recruitment programme before mid-April 2003 and 50 to be recruited over the course of time. Post-doctoral positions at INRIA are intended for students having recently obtained a doctoral degree. The post-doctoral visits are generally of 12 months duration. For detailed information, see: http://www.inria.fr/ travailler/opportunites/ postdoc.en.html

ERCIM - The European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development, in information technology and applied mathematics. Its national member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.

ERCIM is the European Host of the World Wide Web Consortium.

Austrian Association for Research in IT
c/o Österreichische Computer Gesellschaft
Wollzeile 1-3, A-1010 Wien, Austria
Tel: +43 1 512 02 35 0, Fax: +43 1 512 02 35 9
http://www.aarit.at/

Central Laboratory of the Research Councils
Rutherford Appleton Laboratory
Chilton, Didcot, Oxfordshire OX11 0QX, United Kingdom
Tel: +44 1235 82 1900, Fax: +44 1235 44 5385
http://www.cclrc.ac.uk/

Consiglio Nazionale delle Ricerche, ISTI-CNR
Area della Ricerca CNR di Pisa,
Via G. Moruzzi 1, 56124 Pisa, Italy
Tel: +39 050 315 2878, Fax: +39 050 315 2810
http://www.isti.pi.cnr.it/

Czech Research Consortium
for Informatics and Mathematics
FI MU, Botanicka 68a, CZ-602 00 Brno, Czech Republic
Tel: +420 2 688 4669, Fax: +420 2 688 4903
http://www.utia.cas.cz/CRCIM/home.html

Centrum voor Wiskunde en Informatica
Kruislaan 413, NL-1098 SJ Amsterdam,
The Netherlands
Tel: +31 20 592 9333, Fax: +31 20 592 4199
http://www.cwi.nl/

Fonds National de la Recherche
20, Montee de la Petrusse
L-2912 Luxembourg
Tel. +352 26 1925 31, Fax +352 26 1925 35
http:// www.fnr.lu

Foundation for Research and Technology - Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
Tel: +30 2810 39 16 00, Fax: +30 2810 39 16 01
http://www.ics.forth.gr/

Fraunhofer ICT Group
Friedrichstr. 60
10117 Berlin, Germany
Tel: +49 30 726 15 66 0, Fax: ++49 30 726 15 66 19
http://www.iuk.fraunhofer.de

Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
Tel: +33 1 3963 5511, Fax: +33 1 3963 5330
http://www.inria.fr/

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and
Electrical Engineering, N 7491 Trondheim, Norway
Tel: +47 73 59 80 35, Fax: +47 73 59 36 28
http://www.ntnu.no/

Swedish Institute of Computer Science
Box 1263,
SE-164 29 Kista, Sweden
Tel: +46 8 633 1500, Fax: +46 8 751 72 30
http://www.sics.se/

Swiss Association for Research in Information Technology
c/o Prof. Dr Alfred Strohmeier, EPFL-IC-LGL,
CH-1015 Lausanne, Switzerland
Tel +41 21 693 4231, Fax +41 21 693 5079
http://www.sarit.ch/

Slovak Research Consortium for Informatics and
Mathematics, Comenius University, Dept.of Computer
Science, Mlynska Dolina M, SK-84248 Bratislava, Slovakia
Tel: +421 2 654 266 35, Fax: 421 2 654 270 41
http://www.srcim.sk

Magyar Tudományos Akadémia
Számítástechnikai és Automatizálási Kutató Intézete
P.O. Box 63, H-1518 Budapest, Hungary
Tel: +36 1 279 6000, Fax: + 36 1 466 7503
http://www.sztaki.hu/

Trinity College
Department of Computer Science,
Dublin 2, Ireland
Tel: +353 1 608 1765, Fax: 353 1 677 2204
http://www.cs.tcd.ie/ERCIM

Technical Research Centre of Finland
P.O. Box 1200
FIN-02044 VTT, Finland
Tel:+358 9 456 6041, Fax :+358 9 456 6027
http://www.vtt.fi/tte

Order Form

If you wish to subscribe to ERCIM News ***free of charge*** *or if you know of a colleague who would like to receive regular copies of ERCIM News, please fill in this form and we will add you/them to the mailing list.*

Send, fax or email this form to:

ERCIM NEWS
Domaine de Voluceau
Rocquencourt
BP 105
F-78153 Le Chesnay Cedex
Fax: +33 1 3963 5052
E-mail: office@ercim.org

Name:

Organisation/Company:

Address:

Post Code:

City:

Country

E-mail:

Data from this form will be held on a computer database. By giving your email address, you allow ERCIM to send you email

You can also subscribe to ERCIM News and order back copies by filling out the form at the ERCIM website at http://www.ercim.org/publication/Ercim_News/