ERCIM NEWS

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

Number 67, October 2006

Special:

Embedded Intelligence

CONTENTS

Cover illustration by Pertti Jarla

Next issue: January 2007
Special theme: Traffic Planning and Logistics

CoreGRID Summer School 2006 - First Steps in Next Generation Grid Technologies

by Stephan Springstubbe and Wolfgang Ziegler

The CoreGRID Summer School 2006 (CSS'06) took place in Bonn, Germany from 24 to 28 July. Organising the yearly summer school, CoreGRID - The European Research Network on Foundations, Software Infrastructures and Applications for large-scale distributed GRID and Peer-to-Peer Technologies - aims at providing the possibility of advanced training to students and researchers in a trend-setting domain of computer science, which is currently under-represented in the curricula of universities. CoreGRID summer schools are organised by a different CoreGRID partner each year and take place in distinct European cities.

CSS'06 was organised by the Department of Bioinformatics of the Fraunhofer Institute SCAI and hosted by the Bonn-Aachen International Centre for Information Technology (B-IT). The B-IT generously supported the event, providing their whole infrastructure.

The 80 participants included 55 students and researchers from 12 countries and 25 international experts from eight countries, with approximately two out of three of the participants coming from institutions involved in CoreGRID. They worked on 35 course units focusing on Grid Resource Management, Scheduling and Monitoring, Grid Application in Life Sciences and Industry, as well as Future Generation Grids. These course units also included the results of the most important European Grid projects and two half-day hands-on exercises. CSS'06 was financially supported by IBM and SUN, allowing it to maintain a registration fee appropriate for students, who indicated on their evaluation forms that they were very satisfied with the event.

CoreGRID summer school participants.

CSS'06 will have a bearing on several levels, the most important being the curriculum of the B-IT, the further integration of the European Grid research and development, the transfer of knowledge to their home institutions via the students and researchers and, most importantly, the growing support for a broader use of Grid technologies.

For the B-IT the CSS'06 marks the launch of a Grid curriculum. For the first time, starting with the winter semester 2006/07, a course will be offered which is designed as an introduction to Grid technologies and applications. For the following semesters more in-depth courses are planned, in particular addressing the application of Grid technology in the Life Sciences. Researchers of the B-IT expect that Grid technologies will help satisfy the enormous demand for computing power, data and specific services in the Life Sciences and, for example, the pharmaceutical industry, which until now may only be covered using supercomputers. Hence, this aspect of student training has a strategic impact and helps in assuring competitive advantages.

Experience gained during this summer school and from the previous CSS point to the creation of initial networks between participants from different European institutions and the lecturers. Participants and lecturers normally continue maintaining these networks after the event, thus advancing the integration of the European Grid research and development which is - together with excellence - the main objective of CoreGRID.

In the home institutions of the participants, the awareness of research and development around Grid technologies is increasing through the imminent transfer of knowledge by the participants.

Finally, the result of the training of scientists and students across the three levels mentioned helps to support the broader use of Grid technologies in different research disciplines and industry in the future.

The CoreGRID Network of Excellence is administrated by ERCIM.

Links:

CoreGRID: http://www.coregrid.net

2006 Summer School: http://www.coregrid.net/mambo/content/view/231/221/

Please contact:

Stephan Springstubbe
Fraunhofer-Institute for Algorithms and Scientific Computing (SCAI)
Tel: +49 2241 14 2337
E-mail: stephan.springstubbe@scai.fraunhofer.de

Cor Baayen Award 2006 for Oliver Heckmann

Oliver HeckmannOliver Heckmann from Technical University Darmstadt, Germany, and Google Labs Zürich, Switzerland, has been awarded the 2006 Cor Baayen Award for a most promising young researcher in computer science and applied mathematics by ERCIM.

In a particularly tight competition, ERCIM has awarded Oliver Heckmann for the outstanding originality, high impact and quality of his work in the field of network science. To put it in very simple words, Oliver Heckmann tries to give Internet users better quality of service for equal or lower costs. He strives to remove inefficiencies, optimise networks and achieve technological breakthroughs that improve quality of service and/or costs for users. He also sheds light on the trade-off between these two goals. The practical applicability of his results has put him in good contact with many real providers, helped him in acquiring funding, and also led to the recent publication of his book "The Competitive ISP" which also contains many results from his excellent PhD thesis "A System-oriented Approach to Efficiency and Quality of Service for Internet Service Providers".

Oliver Heckmann.

Oliver Heckmann studied applied economics and electrical engineering at TU Darmstadt with award for excellence. He did his PhD in computer science at TU Darmstadt with Prof. Ralf Steinmetz as supervisor and Prof. Jon Crowcroft from University Cambridge, UK as second supervisor. His PhD thesis was awarded the prestigious best dissertation award of the German Computer Science Association for 2004. After his PhD, he was working as an assistant professor / research group head at the Multimedia Communications Lab (KOM), leading a team of 7 PhD student. Recently, he joined Google Labs in Zürich, Switzerland, to continue working on improving and innovating the Internet and Internet services.

Finalists 2006

According to the award rules, each institute was allowed to select up to two finalists from its country. 19 finalists for the 2006 Cor Baayen Award have been nominated by the ERCIM institutes:

- Alexandre Bergel, Ireland
- Rickard Cöster, Sweden
- Ton Dieker, The Netherlands
- Guido Dornhege, Germany
- Olivier Dousse, Switzerland
- Sameh El-Ansary, Sweden
- Serge Fehr, The Netherlands
- Alain Frisch,France
- Tudor Gîrba, Switzerland
- Oliver Heckmann,Germany
- András Kovács, Hungary
- Ivan Laptev,France
- Taneli Mielikainen,Finland
- Harald Øverby, Norway
- Tom Schrijvers, Belgium
- Giuseppe Scarpa, Italy
- Alkis Simitsis, Greece
- Vassilios Solachidis, Greece
- Gem Stapleton, United Kingdom.

The Cor Baayen Award

The Cor Baayen Award is awarded each year to a most promising young researcher in computer science and applied mathematics. The award was created in 1995 to honour Cor Baayen, the first president of ERCIM.

The award consists of a cheque for 5000 euro together with an award certificate. The award winner will be invited to the fall ERCIM meetings.

Cor Baayen Award 2007

Conditions:

- nominees must have carried out their work in one of the 'ERCIM countries': Austria, Belgium, Czech Republic, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Luxembourg, Norway, Spain, Sweden, Switzerland, The Netherlands and the United Kingdom
- nominees must have been awarded their PhD (or equivalent) after 15 April 2004
- a person can only be nominated once for the Cor Baayen Award.

Submitting a nomination:

- nominations should be made by a staff member of the university or research institute where the nominee is undertaking research. Self nominations are not accepted
- to submit a nomination fill out the Cor Baayen Award nomination form at the ERCIM web site (see URL below).

Selection :

- the selection of the Cor Baayen award winner is the responsibility of the ERCIM Executive Committee, who will consult expert opinion in reaching their decision.

Deadline:

- nominations must be submitted by 15 April 2007.

Link:
http://www.ercim.org/activity/cor-baayen.html

Net-WMS - Towards a New Generation of Networked Warehouse Management Systems

NET WMS
Networked Businesses in WMS

by Bruno Le Dantec

Net-WMS is a new project managed by ERCIM with the objective to integrate virtual reality and optimisation techniques in a new generation of networked businesses in Warehouse Management Systems under constraints. The project has been set up in the frame of the ERCIM Working Group 'Constraints'.

Warehouse Management Systems (WMS) are evolving from controlling goods movement and storage to light manufacturing, transportation and order management. Demands from end users are asking for WMS solutions with decision-making applications and networked environment.

Net-WMS project participants.

The Net-WMS project proposes a software solution enabling the expected new generation of networked businesses WMS. Net-WMS will handle networked communication and co-operation processes through the integration of decision-making technologies, generic 3D placement primitives, virtual reality for 3D visualisation, interactivity to design packing models, knowledge modelling.

The added value and competitiveness increase of the Net-WMS solution will be characterised by shared expertise, easy deployment and maintenance, flexibility, interoperability, better resource handling and access to remote services.

The Net-WMS achieved prototype applications for processes in a networked WMS will include:

- a packing modeller of items based on virtual reality and optimisation techniques
- a palletiser tool using optimisation techniques
- a dispatcher with the virtualisation of a truckload
- a set of interfaces will enable Net-WMS communication between several planning components across a network.

Net-WMS-generated knowledge will be shared in networked warehouses, thanks to its interoperability and user friendly interface designed for plant level technician users. Net-WMS optimisation developments will also be made available for the whole community through the enhancement of the Choco open-source system. The achieved optimisation technologies will enable 20% saving for packaging costs. The project gathers nine European organisations having a recognised expertise in logistical and packing problems as well as software development in supply chain optimisation. Net-WMS complementary and multidisciplinary consortium includes 3 SMEs, 2 large manufacturers and 4 academics and research centres.

Net-WMS began in September and has a duration of three years. The project is partly funded by the Sixth Framework Programme of the European Commission. The project consortium is composed of nine partners from research and industry, including the ERCIM members SICS and INRIA.

Link:
http://net-wms.ercim.org

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

Joint DELOS – MUSCLE Summer School on Multimedia Digital Libraries

by Remi Ronchaud

For the first time, the DELOS and MUSCLE Networks of Excellence organised a joint summer school. The event was held in San Vincenzo, Italy from 12 to 17 June 2006 and the topic covered was Multimedia Digital Libraries.

About one hundred European universities and research centres have been involved over the last few years in initiatives addressing multimedia data analysis, computer vision and object recognition, multimedia management in large audio and video collections.

Currently, two Networks of Excellence, MUSCLE and DELOS, are addressing these subjects. The networks are partially funded by the European Commission within the sixth EU Framework Programme, and coordinated by ERCIM.

In the scope of their respective dissemination and training activities, MUSCLE and DELOS organized a joint summer school in June in San Vincenzo, Italy, with the support of CNR Pisa, INRIA, CWI and the University of Modena.

The summer school presented software tools and research strategies to enable users to move away from labor-intensive case-by-case modelling of individual applications, and allow them to take full advantage of generic adaptive and self-learning solutions that need minimal supervision in the field of new generations of digital libraries.

Technologies for multimedia digital libraries span over many disciplines: information systems, knowledge representation, computer vision, audio and image processing, compression and storage, machine learning and information retrieval, multimedia data mining, cross-media analysis, user interfaces and interoperability.

Highly qualified lecturers and international researchers have been invited to address theoretical approaches and practical solutions. Among the invited researchers were:

- James Wang, Pennsylvania State University
- Fernando Pereira, Instituto Superior Tècnico Lisboa
- Nicu Sebe, University of Amsterdam
- Milind Naphade, IBM T J Watson Research Center
- Alex Hauptmann, Carnegie Mellon University
- Padraig Cunningham, Trinity College Dublin
- Jia Li, Pennsylvania State University
- Alan Smeaton – Dublin City University.

This summer school highlighted leading research achievements and improved general knowledge in these fields. It also defined a common vision of future research orientations in multimedia datamining, and provided a significant opportunity to exchange experiences, share views and address common problems. Hopefully, this initiative will stimulate the emergence of new research collaborations.

Links:

Presentations:
http://www-rocq.inria.fr/imedia/DelosMuscleSummerSchool2006/
MUSCLE: http://www.muscle-noe.org/
DELOS: http://www.delos.info/

Please contact:
Remi Ronchaud, ERCIM office
E-mail: remi.ronchaud@ercim.org

Invited lecturer James Wang from Pennsylvania State University.

Summer school participants.

ERCIM Working Group on Software Evolution Spawns Local Chapters

by Tom Mens, Jean-Marie Favre and Salah Sadou

The ERCIM Working Group on Software Evolution has begun to establish some local chapters in order to achieve a better dissemination of results at a national level and to achieve a better interaction with local companies. This is important since the language barrier is still an issue in many countries and the introduction of national activities that are carried out in the native tongue address this problem.

The first national chapter of the Working Group, RIMEL (an acronym for Rétro-Ingénierie, Maintenance et Evolution des Logiciels), was set up in France in spring 2006. Its main objective is to bring together French teams working on reverse engineering, maintenance and software evolution. In France, RIMEL will become part of a national Groupement De Recherche (GDR) on Programming and Software Engineering, an ongoing initiative. RIMEL brings together more than 20 groups from academia and five groups from industry. It promotes software evolution as a major topic in the French research and education community, as well as in industry. In order to tackle the many scientific challenges that arise in this vast and important research domain, RIMEL includes teams from various disciplines, for example database systems, model-driven engineering, software visualisation, and service-oriented architecture to name but a few.

In 2006, RIMEL organised four important events in France in collaboration with the WG on Software Evolution:

- In March 2006, a one-day workshop on software evolution took place, attached to the French-speaking conference Langages et Modèles à Objets (LMO) in Nîmes. The workshop's main goal was to provide PhD students with the opportunity to present their ongoing research in their mother tongue. The workshop style gave plenty of room for lively discussions.
- In April 2006, the main annual workshop of the WG on Software Evolution in Lille was organised in close collaboration with some of the RIMEL members. As such, it attracted 14 French participants, making up more than one third of the total number of participants.
- In June 2006, the official kick-off meeting took place, as part of the French days on Model-Driven Engineering (Ingénierie Dirigée par les Modèles – IDM).
- In September 2006, during the first French-speaking Conference on Software Architectures in Nantes (Conférence sur les Architectures Logicielles - CAL), RIMEL organised two working sessions devoted to architectural evolution and software maintenance and re-engineering, respectively.

RIMEL will continue to be very active in the future, and we hope to exploit the success of this initiative to start similar initiatives in other European countries.

Links:
Working Group web site:
http://w3.umh.ac.be/evol/
RIMEL website: http://planetmde.org/rimel

Please contact:
Jean-Marie Favre, Laboratoire LSR, Institut d'Informatique et Mathématiques Appliquées de Grenoble
Tel: +33 4 76 63 55 71
E-mail: jean-marie.favre@imag.fr

ERCIM Environmental Modelling Workshop

by Thomas Lux

The 13th workshop of the ERCIM Working Group Environmental Modelling was held on 31 August - 1 September 2006 at Fraunhofer FIRST, Berlin, Germany.

Members of the ERCIM Working Group Environmental Modelling participated along with researchers from Fraunhofer FIRST and their co-operation partners in the 13th workshop in order to present their recent research work and to discuss the current trends and developments in the field of environmental modelling.

Research issues that have been addressed include: European operational air quality forecast, satellite data assimilation for air quality forecast, modelling and simulation of processes in the soil and groundwater zone, and especially current aspects of using Grids and Grid workflows for different environmental applications. It was an enjoyable meeting with interesing scientific subjects and fruitful discussions as well as a pleasant social event which gave reason to remember the founding of the working group which was quite exactly ten years ago.

Link:
Working Group web site:
http://wwwold.first.fhg.de/ercim/

Please contact:
Achim Sydow, ERCIM Environmental Modelling Working Group coordinator
E-mail: achim.sydow@first.fraunhofer.de

FMICS 2006 - 11th International Workshop on Formal Methods for Industrial Critical Systems

by Luboš Brim and Martin Leucker

The 11th International Workshop on Formal Methods for Industrial Critical Systems (FMICS 06) was held in Bonn, Germany, on August 26-27, 2006 as a satellite event to the 17th International Conference on Concurrency Theory (CONCUR 2006). The workshop, being also an annual meeting of the ERCIM Working Group, continued successfully the aim of FMICS workshop series - to promote the use of formal methods for industrial applications, by supporting research in this area and its application in industry. The emphasis in these workshops is on the exchange of ideas between researchers and practitioners, in both industry and academia.

Previous workshops were held in Oxford (March 1996), Cesena (July 1997), Amsterdam (May 1998), Trento (July 1999), Berlin (April 2000), Paris (July 2000), Malaga (July 2002), Trondheim (July 2003), Linz (September 2004), and Lisbon (July 2005). The 2006 workshop was organized by the Masaryk University Brno and the Technical University Munich. Forty two registered participants from academia and industry from about eleven countries attended the workshop.

This year the program committee received a record number of submissions. The 16 accepted regular contributions and 2 accepted tool papers, selected out of a total of 47 submissions, cover formal methodologies for handling large state spaces, model based testing, formal description and analysis techniques as well as a range of applications and case studies.

The workshop program included two excellent invited talks, respectively by Anna Slobodova from Intel on 'Challenges for Formal Verification in Industrial Setting' and by Edward A. Lee from University of California at Berkeley on 'Making Concurrency Mainstream'.

The award for the best paper was granted this year to Michael Weber and Moritz Hammer for their excellent paper 'To Store or Not To Store' Reloaded: Reclaiming Memory on Demand" on the techniques for handling extremely large state spaces in model checking of computer systems. The award was granted with the support of the European Association of Software Science and Technology (EASST).

The final proceedings of the workshop will be published jointly with the 5th International workshop on Parallel and Distributed Methods in Verification (PDMC 2006) as post-proceedings in the Springer's Lecture Notes in Computer Science. Selected papers will be invited for publication in a special issue of the International Journal on Software Tools for Technology Transfer. The organizers wish to thank CONCUR for hosting the FMICS 2006 workshop and taking care of many administrative aspects, and ERCIM for its financial support of the workshop. Additionally, the organizers would like to thank EASST (European Association of Software Science and Technology), Faculty of Informatics, Masaryk University Brno and the Technical University Munich for supporting this event.

Best paper award winners Michael Weber and Moritz Hammer.

Link:
FMICS Working Group:
http://www.inrialpes.fr/vasy/fmics/

Please contact:
Pedro Merino Gómez, ERCIM FMICS Working Group coordinator, Universidad de Málaga / SpaRCIM
Tel: +34 952 132752
E-mail: pedro@lcc.uma.es
http://www.lcc.uma.es/~pedro

A Security and Dependability Strategic Research Agenda for Europe

by Jim Clarke

Europe's research framework programmes are committed to the establishment of a solid security and dependability infrastructure. The IST-SecurIST project has been charged with the preparation of a European strategic research agenda in the field of ICT for Security and Dependability, for the upcoming 7th research framework programme.

The Internet and other digital networks have now become an integral part of our economy and society. But as we are rapidly adopting more information and communication technologies (ICT) in services and commerce, private information is at increasing risk and security and dependability problems become prevalent. Indeed, today people are becoming more and more concerned about the increasing complexity of information and communication systems and the proliferation of privacy-invasive information gathering sources and techniques. In their online daily interactions, they often find themselves faced with high-profile losses of their personal information and with viruses, spam, phishing and other crimes of growing severity and sophistication. As a result, they find themselves in an undesirable situation in which they must put ever more trust into environments in which there is little or no way of understanding or assessing them properly.

To build an information society that will deliver growth and prosperity, we need to tailor ICTs to business and social needs, and ensure that they become useful tools for economic and social innovation. The starting point for making them useful is to foster trust and safeguard security in a networked world. In this respect, Europe's research framework programmes are committed to the establishment of a solid security and dependability infrastructure. The IST-SecurIST project (an IST Coordination Action) has been charged with the preparation of a European strategic research agenda in the field of ICT for Security and Dependability, for the upcoming 7th research framework programme (FP7, 2007–2013). In order to achieve this objective, the SecurIST project has established two fundamental bodies: the European Security and Dependability Task Force (STF), and the SecurIST Advisory Board.

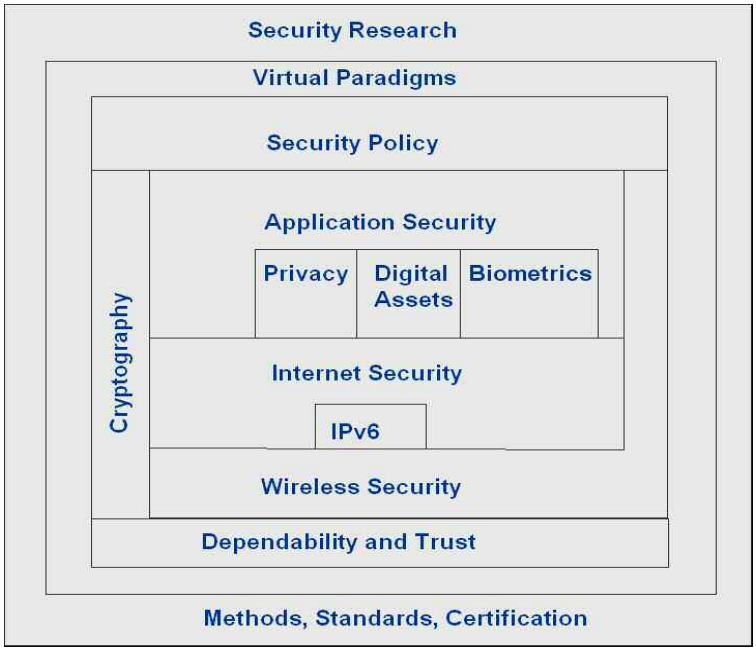

Security and Dependability Task Force initiatives integration.

The STF is currently comprised of 180 members spread across thirteen fundamental thematic areas (initiatives) of research. It provides a forum for consolidation and consensus building. The thematic initiatives are shown in the figure, which provides a visual interpretation of how these initiatives are integrated and work together.

The SecurIST Advisory Board is composed of European experts in information security and dependability. The charter of the board is to oversee, review, enhance and promote results from the STF (see http://www.securitytaskforce.eu).

Recently, based on inputs from the STF, the SecurIST Advisory Board has issued a document presenting its recommendations for a future security and dependability research framework in Europe, for the period 2007-2013. Under the headline From 'Security and Dependability by Central Command and Control' to 'Security and Dependability by Empowerment', the Advisory Board is recommending the following nine key research areas:

1. *Empowerment of the stakeholders:* Stakeholders of the information society include individual citizens, industry and academia, non-governmental organisations and governments. Empowerment of the stakeholder is vital as there is a clear technological trend towards decentralization of technology, as well as of its management and control. Responsibility, authority and control have to move more towards the end user.

2. *Europe-specific security and dependability:* Europe has a very specific heterogeneous culture and history and set of attitudes to trust and society that requires specific research profiling.
3. *Infrastructure robustness and availability:* Further research efforts are needed for the assurance of ICT network and service infrastructures, as well as the robustness and availability of critical infrastructure, such as health, energy, transport or finance.
4. *Interoperability:* Research on the interoperability between security and dependability technologies and standards.
5. *Processes for developing secure and dependable systems:* Research on the systematic improvement of secure and dependable system development (including hardware and software) from their design phase.
6. *Security and dependability preservation:* In an increasingly complex world of evolving requirements, technologies and systems, maintenance of effective system security and dependability is critical and is essential for preserving user confidence.
7. *User-centric security and dependability standardisation:* Strengthen the structured involvement of end users and their respective representatives into relevant standardization activities involving security and dependability technologies.
8. *Security and dependability of service oriented architectures (SOA):* The need to establish and maintain trust and manage policy regulations and service level agreements in an SOA context, together with commensurate advances in software engineering to deliver service expectations.
9. *Technologies for security:* Research is needed to provide a higher assurance of trusted communication and handling of digital information. Cryptology and trusted functionality and computing need to be considered.

In addition to these nine key research areas, four future grand challenges covering a 10-20 year vision are presented by the Advisory Board. They illustrate possible longer-term possibilities and implications:

1. *Countering vulnerabilities and threats within digital urbanization:* This challenge addresses open problems that we will face in security and dependability from the expansion and globalization of digital convergence by 2010-2015.
2. *Duality between digital privacy and collective security:* digital dignity and sovereignty: This challenge deals with future privacy issues of all the stakeholders, whether citizens, groups, enterprises or states. It addresses the problem of how to override the 'Big Brother' syndrome and 'dark security', ie, the future assurance of digital sovereignty and dignity for the various stakeholders.
3. *Objective and automated processes:* This challenge addresses the problem of how to attain a controllable and manageable world of complex digital artefacts by 2015 and how to inject regular, quantitative techniques and engineering to make the field truly scientific.
4. *Beyond the horizon: a new convergence:* This last challenge deals with the preparation of a new convergence at a horizon of 2020 and beyond, which is the bio-nano-info-quantum 'galaxy' and the new security and dependability challenges that will emerge.

Links:
The full report of the SecurIST Advisory Board recommendations:
http://www.securitytaskforce.eu/

Please contact:
Jim Clarke, Waterford Institute of Technology, Telecommunications Software and Systems Group, SecurISTCoordinator
E-mail: jclarke@tssg.org

European Research Advances in Embedded Intelligence

The IST Results news service features numerous European research organisations and industrial partners whose work on embedded intelligence has been funding under the European Commission's Information Society Technologies (IST) R&D priority.

These research projects focus on many different applications of embedded systems, for example:

- the ASAP project team has produced an open source programming, analysis and optimisation toolkit for pervasive computing systems using Constraint Logic Programming (CLP) languages. It has been validated in a series of case studies and is expected to open the way to more effective and efficient embedded systems development in the future.
- The Embedded WiseNts project has brought together twelve partners from the top research institutions in wireless communication, distributed computing and cooperating objects in ten different European countries to focus on the development of Wireless Sensor Networks (WSN) and their applications, especially in the form of Cooperating Objects (CO), to help develop a roadmap for innovative future applications.
- A new technology, dubbed Embedded and Communicating Agents, has allowed researchers at Sony's Computer Science Laboratory in France to add a new level of intelligence to a robotic dog. Instead of teaching the dog new tricks, the algorithms, design principles and mechanisms developed by the ECAgents project allow the robotic pet to learn new tricks itself and share its knowledge with others.

For more examples of recent research results related to embedded intelligence and other current ICT themes, visit IST Results at http://istresults.cordis.europa.eu

W3C Workshop on Mobile Web in Developing Countries

W3C invites experts to participate in the 'Mobile Web in Developing Countries' Workshop in Bangalore, India, on 5-6 December 2006.

Participants will discuss the challenges, requirements, and use cases for mobile Web access in developing countries and thus help bridge the digital divide. The Workshop will bring together experts in mobile Web technologies and specialists on emerging countries and the digital divide.

This public Workshop is part of W3C's Mobile Web Initiative, which aims to identify and resolve challenges and issues of accessing the Web when on the move.

W3C invites support for this Workshop through a three-tier sponsorship program designed to support participation by people or organizations who might otherwise not have the financial means to attend the meeting.

Sponsors reach those who are making decisions about the future of the Web, as well as those in the public who have come to rely on the Web as critical infrastructure for development.

To participate in the Workshop, please submit a position paper by email before 1 November 2006.

Links:
CfP: http://www.w3.org/2006/07/MWI-EC/cfp.html
Sponsorship Program:
http://www.w3.org/2006/07/MWI-EC/sponsors.html

W3C Launched Secure Browsing Initiative

Recognizing the challenges people face when browsing the Web, W3C launched an initiative to build a foundation for a more secure Web. The new Web Security Context Working Group will propose standards that will enable browsers to do a much better job helping people make proper trust decisions.

Part of W3C's Security Activity, the group's mission is three-fold: to build consensus around what information people need from browsers in order to understand their "security context," to find innovative ways to present this information and raise awareness, and to suggest ways to make browsers less susceptible to spoofing of user interfaces that are used to convey critical security information to end users.

W3C chartered this new work after a successful Workshop on Usability and Transparency of Web Authentication in March 2006. Per the charter, the group will conduct its technical work in public and will operate under the W3C Royalty-Free Patent Policy. Organizations interested in security over the Web should join W3C to participate in this new Web Security Context group.

Links:
Web Security Context Working Group: http://www.w3.org/2006/WSC/
W3C Security Activity: http://www.w3.org/Security/
Join W3C: http://www.w3.org/Consortium/membership

W3C Workshop on XSL-FO 2.0 Requirements

W3C is holding a Workshop to gather inspiration, needs and techniques for a future version of XSL-FO, the formatting specification on 18 October 2006 in Heidelberg, Germany. The W3C XSL Working Group expects that the enhancements for XSL-FO 2.0 will focus on layout-driven formatting, augmenting the content-driven layout facilities already defined. The XSL Working Group is looking for experts in the field to bring their knowledge, insight and experience, as well as for use cases and representatives of user communities.

Link:
http://www.w3.org/Style/XSL/2006-Workshop/

Latest W3C Recommendations

- XHTML-Print
  20 September 2006, Jim Bigelow, Melinda Grant
- Extensible Markup Language (XML) 1.0 (Fourth Edition)
  16 August 2006, Jean Paoli, C. M. Sperberg-McQueen, Tim Bray, François Yergeau, Eve Maler
- Namespaces in XML 1.0 (Second Edition)
  16 August 2006, Dave Hollander, Tim Bray, Richard Tobin, Andrew Layman
- Namespaces in XML 1.1 (Second Edition)
  16 August 2006, Dave Hollander, Richard Tobin, Andrew Layman, Tim Bray
- Extensible Markup Language (XML) 1.1 (Second Edition)
  16 August 2006, Eve Maler, Jean Paoli, C. M. Sperberg-McQueen, Tim Bray, François Yergeau, John Cowan

A complete list of all W3C Technical Reports:
http://www.w3.org/TR/

COPRAS 'ICT Research and Standardization' Conference

Brussels, 17 January 2007

The Cooperation Platform for Research and Standardization (COPRAS), together with its consortium partners CEN, CENELEC, ETSI, W3C and The Open Group, will organize a major conference on ICT research and standardization, in the perspective of FP7 on Wednesday 17 January 2007.

The conference will bring together all stakeholders involved in the research/standards interfacing process, and will feature well known speakers from both the research and the ICT standards communities, as well as representatives from the European Commission. Areas to be addressed include:

- how passing research output through standards processes supports innovation, and helps projects bringing their results to the market
- what are the major challenges IST projects face when cooperating with standards organizations, and how these challenges should be addressed
- how should interfacing between research and standardization be organized in FP7, and what are the role and challenges for the European Technology Platforms in this respect
- what are the main areas in ICT standardization where future closer cooperation with research projects in FP7 will be required
- how can contributing to standards-making be made more attractive to research projects, and how this will contribute to Europe achieving itsoverall goal of furthering the Information Society and its position in ICT development
- what are the recommendations from FP6 for the research/standards interfacing process in FP7, and how can COPRAS' results support improving cooperation between projects and standards organizations in the new Framework Programme.

The conference will take place at the Bedford Hotel in Brussels, and participation will be free of charge on a first-come/first-served basis. Pre-registration for the conference is possible at: http://www.w3.org/2004/copras/meetings/open07/registration.html

Links:
COPRAS Project: http://www.copras.org/
COPRAS contact: info@copras.org

W3C Mobile Web Seminar in Paris

Paris, 16 November 2006

W3C announces a public seminar on the technical solutions for the special challenges of mobile Web adoption. The event will be held in central Paris, on Thursday 16 November 2006.

"People could access the Web at all times and in all situations, thanks to mobile devices, but Web sites aren't always ready for the diverse range of devices," explains Philipp Hoschka, Mobile Web Initiative Leader. "W3C's Mobile Web Initiative has already produced concrete results to help developers make their Web sites accessible to mobile devices. This seminar gives a valuable opportunity to discuss those results and state of the art techniques."

W3C's Mobile Web seminar is a public event which will focus on current results produced by W3C's Mobile Web Initiative (MWI). The confirmed speakers include representatives from MWI sponsors such as France Telecom, Jataayu Software, MobileAware, Opera Software, and Vodafone.

This seminar is funded by the Web Beyond 3G (3GWeb) project, financed by the European Commission's FP6 IST programme. Attendance to the seminar is free and and open to the public, but registration is required.

Links:
Mobile Web Seminar: http://www.w3.org/2006/11/mwi-seminar.html
Registration: http://www.w3.org/2002/09/wbs/1/mwiparis/

First W3C Webinar on Mobile Web Design

W3C organized a W3C webinar on 26 July 2006, where Cameron Moll, an expert in mobile Web design, spoke about mobilizing Web content. This presentation took a look at the current state of the art, and how to best extend an existing or planned Web site to handheld devices with minimal repurposing of code, content, and images. It also showed which elements of forward thinking design were needed to maintain consistency, regardless of where and how users access your content. As part of the 3GWeb project, W3C plans to organize more webinars to be announced on the Mobile Web Initiative home page.

Link:
W3C Mobile Web Initiative: http://www.w3.org/Mobile/

Introduction to the Special Theme

Embedded Intelligence

Illustration by Pertti Jarla.

by Erwin Schoitsch
and Amund Skavhaug

Embedded Systems are a key if Europe is to remain in the forefront of digital technology and as such they have been classified as an important research area for the European Union's Seventh Framework Programme - the main financial tool through which the EU supports research and development activities. The IST/FET (Future and Emerging Technologies) programme 'Beyond the Horizon', coordinated by ERCIM, points out that Embedded Systems, in combination with pervasive - or ubiquitous - computing, (cognitive) intelligence and software-intensive systems, which in fact means 'embedded intelligence', or 'smart systems' in the broader context of 'smart environments', are the most important challenge for strategic, long-term research, with a huge impact on society and the economy. The ITEA2 Roadmap (Information Technology for European Advancement) reaches the same conclusion - that Embedded Systems are a crucial technology for European competitiveness.

ARTEMIS (Advanced Research and Technology for Embedded Intelligence and Systems) is a strong, industry-driven European Technology Platform (ETP) which aims to establish a coherent, integrated European research and development strategy for embedded systems (http://www.artemis-office.org).

EPoSS, another ETP launched in July this year (see separate article by the author in this edition) focuses on smart systems integration, which is considered an important emerging area. The key aspects are building systems from components, interdisciplinarity, a holistic approach to pervasive and ubiquitous computing, fast integration of a variety of technologies, energy autonomy and networking (http://www.smart-systems-integration.org).

Several national research programmes in Europe cover essential aspects of this theme, for example FIT-IT in Austria (BMVIT, Federal Ministry for Transport, Innovation and Technology), with topics such as Embedded Systems, System on Chip, Semantic Systems and Security, focusing on radical innovations in these areas.

'Intelligence' takes account of autonomous reasoning and acting in a co-operative manner. 'Ambient Intelligence' refers to an exciting new paradigm in information technology, "in which people are empowered through a digital environment that is aware of their presence and context and is sensitive, adaptive and responsive to their needs, habits, gestures and emotions." (ISTAG scenarios, and Schoitsch, E., Bloomfield, R. et al. (2003), AMSD – Dependable Embedded Systems Roadmap, IST project 37553) . This applies not only for people-centred tasks, which, of course, seems the most exciting, science-fiction-type, aspect, but also for purely technical solutions like smart sensors, actuators and control systems, especially in safety related applications. Heterogeneity (of environment, applications, protocols, etc.), autonomy (self-awareness, self healing, self-organizing, etc.), nomadic mobility (ad hoc, unreliable, heterogeneous, etc.) and scale-less (number of users, geography, structure, etc.) are the new emerging embedded systems challenges (Neeraj Suri, Keynote at DECOS Conference, ME'06 Conference, 2006).

This special theme fits in very well with the current European framework and strategic research discussions. The areas addressed - with the related European research projects are referenced in brackets - include:

Embedded Systems Applications

This part provides an overview of four different, but nevertheless typical applications of embedded systems, ie, some work on autonomous systems, on intelligent road safety through co-operative networks (COOPERS); on ambient assisted living support for the 'aging society' (BelAmI) and on an experiment with iCat, a personal robot platform with emotional feedback. This wide spectrum of projects is typical for the broad range of foci addressed by the engineering of AmI applications for the future intelligent environment of people.

Embedded Systems Design and Development

This provides an overview of the DECOS (EU-FP6 project) tool chain for the development of dependable embedded systems in a cost-efficient manner; of a project on reconfigurable hardware for 'Embedded Architectures on Demand' (AHEAD); on system-level design of fault-tolerant systems (SynDEx); on a tool to model the behaviour of embedded system designs (mCRL2); on reactive processing for reactive embedded sys-

tems and on a very specialised aspect of designing concurrent systems (tagged procedure calls, TPC). This covers a broad range of system architecture, design and development issues in the embedded systems area.

Ambient Intelligence (AmI) Systems, Smart Environments, Platforms and Services

The articles provide an overview of one ERCIM member's AmI research facility (FORTH); on a user-centred design approach to smart environments (EASE); on platforms for ambient services - spanning the gap between the issue of sensor networks and adaptive applications - (Ambiance, Construct); on Service-Oriented Architectures for pervasive computing environments (ARLES) and, as a special case, intelligent materials for smart applications. This part should give a feeling for the broad range of perspectives of 'ambiance' (vs. embedded systems in technical engineering applications).

Further on, this special theme continues with articles covering specific key aspects and technologies of embedded systems and ambiance implementations:

Wireless Sensor Networks

Firstly, there is a description of a Network of Excellence to overcome the fragmentation of research in this field (CRUISE). Then two key aspects of wireless sensor networks are highlighted: Intelligent RFID sensor integration (IntelliSense RFID) and collaborative capture, linking information systems to real work by sensor networks (ACES).

Embedded Systems

A key characteristic of embedded systems is the use of restricted resources, a problem to be resolved beforehand by modelling. One article covers autarkic power generation for networked embedded systems (FIT-IT, Intermon), another covers guarantees on resource usage bounds (EmBounded).

Massively Deployed Embedded Systems: Trust and Security

Massively deployed and networked embedded systems are prone to security breaches leading to loss of safety or reliability/availability of safety or loss of business critical services (and loss of consumer confidence). The ITEA project Trust4All project has written two articles on different aspects of this problem - research on context-aware trust in open environments as a basic research issue and a trust management framework. The special problem of protecting privacy by providing separate digital identities depending on the context of use in ambient environments is handled in another article (BASIS).

These areas overlap in part, so overlapping topics may be addressed as well. This theme should provide a vision of the future for anyone interested in embedded systems and address the challenges and risks for research, the economy and society rather than focusing on a single, isolated subtopic. In most cases, the articles are based on reports of European research projects and Networks of Excellence and thus provide a good overview of research and applications in the areas addressed.

Several ERCIM working groups and many teams within ERCIM member institutes are involved in at least some of the aspects of this special theme, which was suggested by the ERCIM Working group on Dependable Embedded Systems. We have tried to provide a comprehensive overview on many different aspects of this important theme.

Please contact:
Erwin Schoitsch, ARC Seibersdorf research/AARIT, Austria
E-mail: erwin.schoitsch@arcs.ac.at

ARTICLES IN THIS SECTION

Autonomous Systems - Safety Critical Embedded Systems and Intelligence

by Erwin Schoitsch and Wilfried Kubinger

Embedded systems are everywhere today and will be even more prevalent in the future. Most of them perform simple control and support tasks on command in a pre-programmed manner, but that's not the end of the line. Increasingly, they are designed to carry out autonomous tasks including 'intelligence' (smart systems). They are also designed to be able to decide, based on complex inputs and situation awareness, what to do under circumstances defined by an unreliable environment. Autonomous vehicles are the next step in this direction.

Basic characteristics of embedded systems are that distributed, networked electronic control units (ECUs) are integrated into their technical or natural environment, linked to sensors and actuators with properties like cognition and situation awareness and controlled by robust algorithms which implies some 'intelligence' (smart systems). Applications are to be found in many domains: safety-critical applications include aerospace, railways, automotive applications, machinery and medical systems, as well as entertainment. Challenges relate to the dependability (safety, reliability and security) of hardware, software and systems design, communication technologies, sensors, actors, materials etc. The Integrated Project DECOS is carrying out research into the design and development of holistic dependable systems, providing the fundamental architecture and services in a domain and platform independent manner.

The next step on from a collection of 'intelligent embedded functions' is the development of autonomous systems. These are able to perform complex missions in an autonomous manner, coping with unexpected incidents and interactions from the environment. They need to meet two basic criteria in the range of their activity, especially in case of interaction with people:

- a dependable, robust infrastructure as a basis of the system (as developed, for example, by DECOS)
- robust and reliable algorithms, sensors and actuators, situation awareness and cognition.

The next step would be co-operative systems consisting of an ad hoc set of autonomous systems (eg autonomous platooning of cars on a highway, controlling the traffic-flow by car-to-car communication).

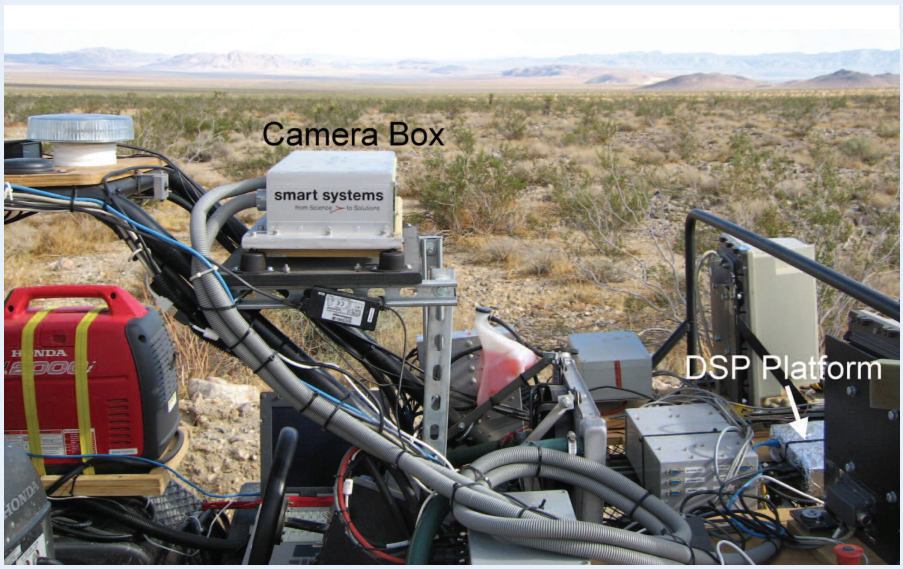

Figure 1: RASCAL on track at DARPA Grand Challenge '05 (with stereo vision sensor).

ARC Seibersdorf research is doing research in the area of autonomous systems, especially in the area of automotive and other mobile transportation systems, providing vision systems (optical and laser light systems), advanced driver assistance systems, recognition/identification of, for example, traffic signs, demonstrations of platooning systems (model cars) and with robots (partner in the recently accepted robots@home proposal, a STREP of FP6).

DARPA Grand Challenge, Automotive Applications

A significant milestone in the development of autonomous vehicles was the DARPA Grand Challenge (Defense Advanced Research Projects Agency, USA). The driverless vehicles had to complete a course over 132 miles (212 km), through desert land with off-road type terrain in the Mojave Desert, without any external support.

Figure 1 shows SciAutonics/Auburn Engineering's autonomous vehicle, RASCAL (Robust Autonomous Sensor Controlled All-Terrain Land Vehicle). They invited ARC Seibersdorf research's Smart-Systems team to provide their embedded stereo vision module SMART EYE SVS (including all the algorithms for detection and identification of obstacles, lanes, tracks, etc.). The vision system performed beyond expectations: RASCAL was one of the fastest vehicles in the qualification event, where only 23 out of 40 (of more than 160 that applied) qualified for the final race. RASCAL came 16th in the end, because of a failing communication hub.

The 2007 DARPA Grand Challenge will take place in an urban environment. This means that if the mission cannot be fulfilled because of a closed lane, the vehicle has to turn and autonomously

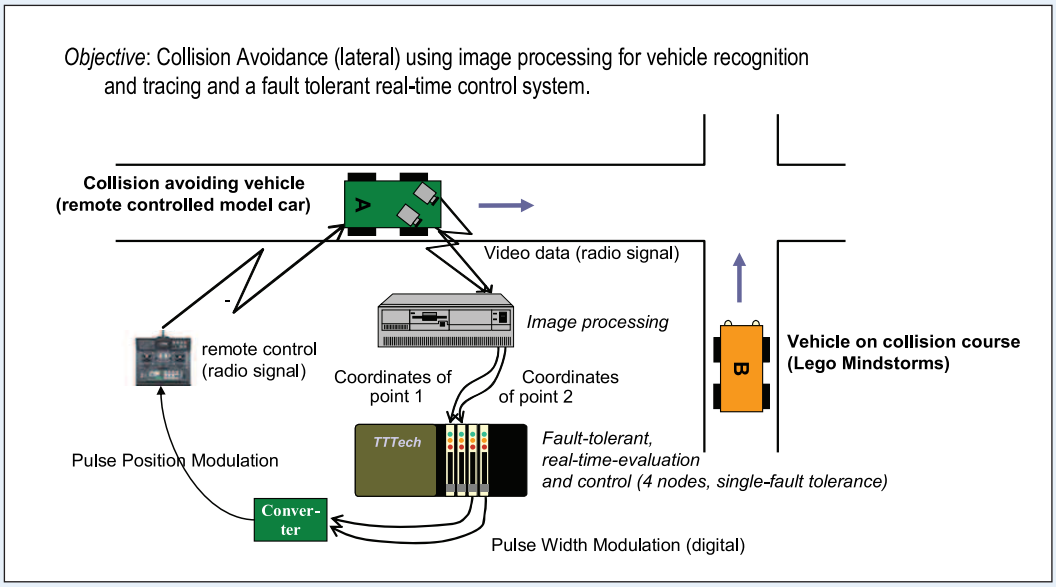

Figure 2: TTIP - Time-Triggered Image Processing Prototype.

decide which alternative route to take, obeying traffic rules.

The embedded stereo vision sensor will be tested in other environments too (a concept car for advanced driver assistance systems) and in a platooning experiment with model cars.

Figure 2 shows another demonstration of highly-dependable system technology – the TT Vision node application (Time-Triggered Vision node, applying TT technology to a demonstrator, avoiding a crash by either stopping or speeding up depending on the situation).

Autonomous vehicles are of primary importance in dangerous environments where human drivers would be at risk of their lives and will, in future, allow safer transport on roads.

Autonomous Robots

Other short-distance vision systems in industrial environments for autonomous robots have been tested with laser light systems. Figure 3 shows such a laser light stereo system.

The advantage of the laser light section system, which was implemented in a feasibility study for a mobile robot system in an industrial environment, is its simplicity and low price.

The smart systems team of ARC Seibersdorf research is a partner in the STREP initiative 'robots@home' of the last IST Call 'Advanced Robotics'. The goal is to develop an open mobile platform that is able to navigate in the different home environments encountered in European and international homes. The home environment poses another set of challenges yet to be resolved: changing environment, tables and chairs with thin feet, glass objects, humans crossing the way, different materials and objects difficult to identify, unreliable environment. The robot needs to be able to 'learn' a new environment in a natural manner by accompanying a person around a room, it needs to be able to classify objects and obstacles and react to different stimuli - and all without the owners doing some type of 'programming' or 'configuration'. This open platform is intended to be a building block for service and personal robots for many different home applications and ambient assisted living tasks.

The consortium is composed of teams from TU Vienna, ARC Seibersdorf research (Austria), ETH Zurich, BlueBotics, Securitas SA, Nestle Nespresso SA (Switzerland), Legrange SA (France), and Otto Bock HealthCare GmbH (Germany).

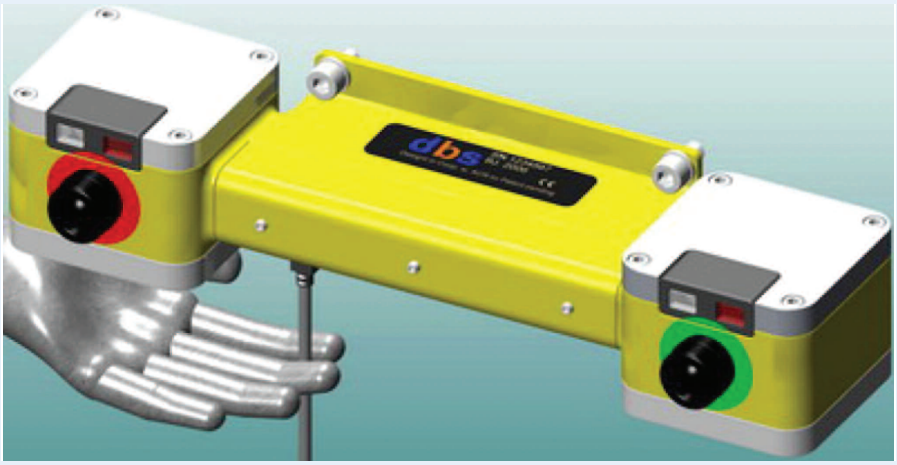

Figure 3: Laser Light Section System, example for mobile robot application.

Links:

http://www.smart-systems.at

http://www.decos.at

Please contact:

Wilfried Kubinger, Erwin Schoitsch, ARC Seibersdorf research/AARIT, Austria

Tel: +43 50550 4200, 4117

E-mail: {wilfried.kubinger, Erwin.schoitsch}@arcs.ac.at

Automotive Visions beyond In-Car Driver Assistance: Integrated Traffic Management with Coopers

by Thomas Gruber and Erwin Schoitsch

Vehicles connected via continuous wireless communication with road infrastructure on motorways, exchanging data and information relevant for the specific road segment to increase overall road safety and enable co-operative traffic management – this is the vision of the Coopers project.

In the Sixth Framework Programme of the European Commission, one of the main themes deals with road traffic safety. Several projects funded by the Sixth IST Framework Programme address this topic, but Coopers takes a specific position with unique ways of achieving a safety improvement through an intelligent network which exploits existing technologies for co-operative services.

Some 38 partners from 15 European countries with industry, university and research institution backgrounds formed a consortium under the leadership of Austria Tech, a company founded by the Austrian Ministry for Transport, Innovation and Technology for driving technological development. The kick-off meeting took place in early spring 2006 and since then the innovative work has been driven by the partners in the areas of system and service architecture, safety analysis and assessment and evaluation of technologies.

Today, there exists a limited set of measures for improving road safety. Those responsible introduce legal regulations and mainly static roadside installations which warn drivers of bends, humps, narrowings and other static danger. But when it comes to accidents or traffic congestions, at present there is little guidance for the driver. Road traffic authorities and road operators rely mainly on the driver's attention, experience and capabilities. Radio traffic news is often late and inaccurate in respect of the location of the event in question. VMS (variable message signs), which could give concrete and more precise advice in case of situation-related dangers, are expensive and therefore rarely used. Currently, the high cost for VMS deployment prevents European road operators adopting this equipment for major road sections, so missing out on a huge safety gain.

Coopers is preparing the way for improving road safety on motorways at an affordable cost. Based on existing technologies and infrastructure, the driver is provided with real-time data on the current traffic situation ahead of him.

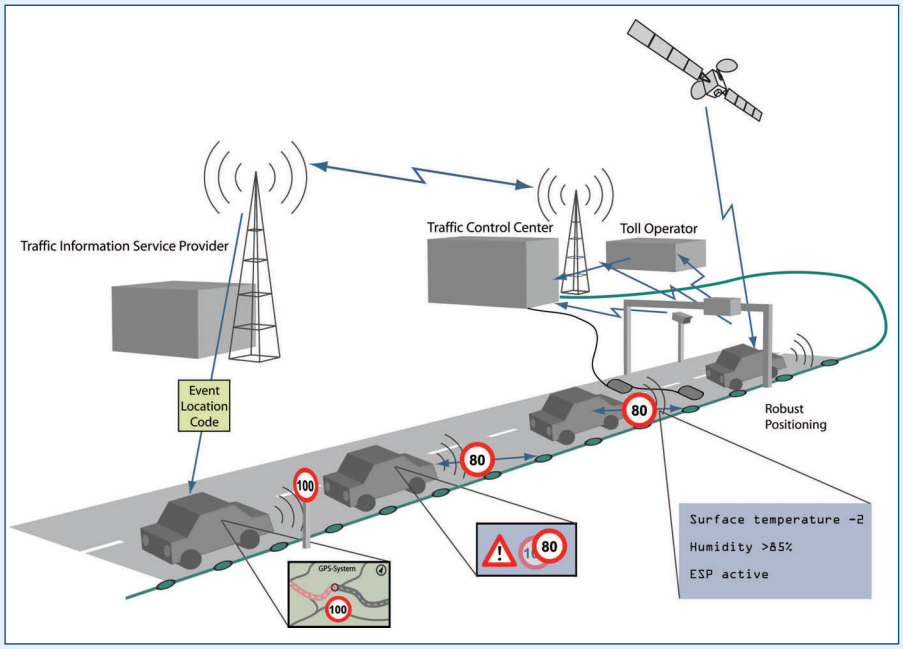

Intelligent Infrastructure and Smart Cars plus individual location based services – I2V and V2I communication.

Coopers Mission:

To define, develop and test new safety related services, equipment and applications using two way communication between road infrastructure and vehicles from a traffic management perspective. COOPERS will build upon existing equipment and infrastructure as far as possible to incorporate bi-directional infrastructure-vehicle links as an open standardised wireless communication technology. The role of motorway-operators in offering and retrieving safety relevant and traffic management information for specific road segments on European motorways based on infrastructure and in-vehicle data will be investigated.

Work Areas in Coopers

- Safety related traffic management and information services
- Roadside and vehicle based data acquisition
- Traffic Control Centre – TCC Applications
- Bi-directional I2V / V2I link
- Roadside Transmitter
- On Board Unit extension and integration
- Standardisation.

Each car is equipped with a receiver for the I2V (infrastructure to vehicle) communication encapsulated in an OBU (onboard unit) and a display gives information about accidents, traffic jams, road construction sites and other location and time related events. Only messages relevant to the driver on a particular segment are passed on to him. There is no irrelevant data about traffic congestions or accidents in remote areas of the country as happens with traffic news on the radio. The information is accurate and precise both in terms of location and time.

Because the driver will be depending on the reliability and accuracy of information, I2V will strongly increase the responsibility and liability of the infrastructure operator. The most effective use of I2V communications is expected in areas of dense traffic where the accident risk and the probability of traffic jams are very high. In addition to V2I and in parallel to it, but in the reverse direction, ie by a I2V (infrastructure to vehicle) link, the real time communication can also be used for verifying infrastructure sensor data using vehicles as floating sensors.

Coopers has started with an assessment of existing technologies, of possible safety improvements and a selection of those services most appropriate for testing the new approach. The work plan comprises three steps for implementing the I2V communication:

- improve road sensor infrastructure and traffic control applications for more accurate traffic information with a infrastructure-to-vehicle communication concept utilising a link to road tolling systems
- develop adequate communication concepts and applications ensuring reliability, real-time capability and robustness considering different technologies like DAB or CALM.
- demonstrate results on important European motorway sections with high traffic density (Rotterdam-Antwerp, Berlin/Darmstadt, Nuremberg-Brenner-Verona) and increase development strategies.

Coopers is compliant with the objective of sustainable transport and development defined in the EC White Paper on Transport policy. As a result it will demonstrate a prototype for future intelligent traffic management systems.

Links:
http://www.smart-systems.at
http://www.coopers-ip.eu/

Please contact:
Thomas Gruber, Erwin Schoitsch, ARC Seibersdorf research, Austria
E-mail: {Thomas.Gruber, Erwin.Schoitsch}@arcs.ac.at

Embedded Intelligence for Ambient Assisted Living

by Frank Bomarius, Martin Becker and Thomas Kleinberger

Our aging societies are facing the problem that more and more elderly people want to live longer at home, in their preferred environment, despite diseases and handicaps. We have a great opportunity to tackle this major demographical problem by building embedded intelligence into systems for Ambient Assisted Living (AAL).

Demographical and social changes have an enormous effect on health care, emergency and welfare services. As the average age continues to rise, with a corresponding increase in chronic diseases, there needs to be a dramatic growth in assistance and care, resulting in even higher service costs, a decrease in quality of service, or both. While today's commercially available emergency and care equipment and health and fitness devices already cover a broad range of features, they are mostly stand-alone systems and are often just too difficult for the elderly to operate without assistance.

Our ambient intelligent care and assistance (amiCA) system will increase quality of life by using small, unobtrusive sensors which communicate wirelessly and are energy saving. We employ sensors that can be easily installed in elderly people's homes and need not be attached to the body, where possible. amiCA will reason on raw data series from individual sensors as well as on combined sensor data to derive more accurate, reliable and/or more abstract data.

By tracking the data, the daily routine of the elderly can be learned and monitored. Significant deviation from the typical routine is used as a predictor of potentially unhealthy behaviour, for example forgetting to drink or eating spoiled food due to dementia, or of dangerous situations such as being too weak to get out of bed and call for help.

By tracking vital parameters, such as indicators of dehydration, of an impending heart attack, of hypoglycaemia, or a dangerously low body temperature, potential health-induced emergency situations can be detected early on. In this way we should be able to avoid or mitigate chronic health impairment as a consequence of delayed treatment. This will effectively improve quality of life for the elderly and cut the costs of health care by shortening or avoiding hospitalisation and subsequent rehabilitation measures.

Following a staged model, amiCA will call for early assistance from family, friends, or care givers. It will provide doctors with tracks of vital parameters and thus improve the precision of diagnosis and treatment. In case of an acute emergency, such as a fall or heart attack, an automatic call to an emergency call centre will be triggered without intervention from the patient, who may not even be able to make the call.

Approach

We have identified major problem areas for the elderly by talking to care personnel, geriatricians, medical doctors and emergency centre staff and have then determined appropriate indicators that can be measured by sensors. Based on this, we set up our Ambient Assisted Living Demonstrator Laboratory (AAL Lab), which aims to come up with embedded intelligence solutions for elderly people living alone at home.

The AAL Lab will be continuously extended. In the first version, to be opened to the public in October 2006, the following scenarios are just a few of those which will be implemented:

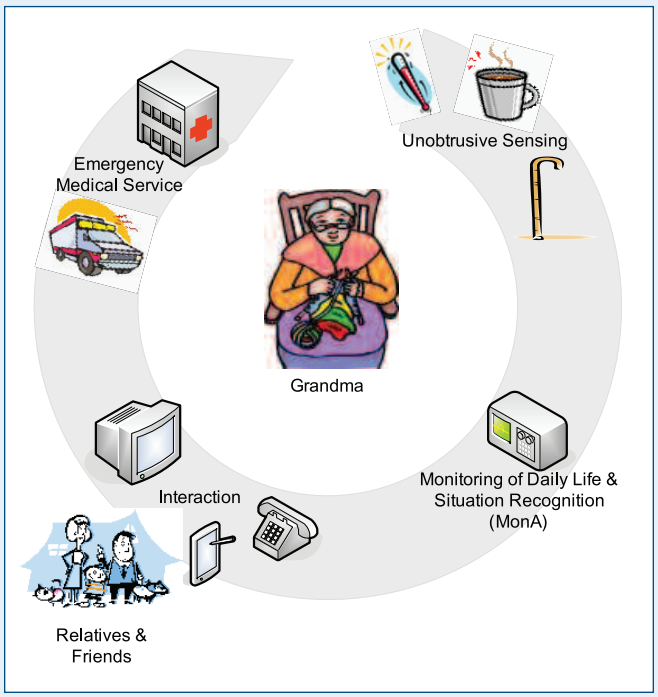

Ambient Intelligent Care and Assistance (amiCA).

Monitored Drinking – Dementia is a major reason for moving the elderly into nursing homes. Dehydration exacerbates dementia and other health problems. Reminding the elderly to drink enough is an important function. Our computerised cup measures the amount of fluids consumed each day by the person handling the cup. Potential dehydration can thus be detected.

Monitoring Food Quality – Eating spoiled food is another source of health problems. An RFID-based system built into a refrigerator checks for expired food and issues warnings. Food related incidents can be prevented or can be traced back to the root cause, allowing for a better treatment.

Location Tracking – Knowing the exact location of a person enables a whole range of location-aware services in the AAL environment, such as switching the lights on or off or activating those communication devices which are closest to the person. RFID labels embedded in the carpet of the AAL Lab are read by a walking aid, allowing the location of the person to be determined.

Fall Detection – A sudden fall is a strong indicator of an emergency situation and is often detected too late, sometimes hours or days later. Gyro sensors built into wrist badges, walking aids, or other body-mounted devices can detect sudden falls by a person and trigger a staged emergency reaction of the AAL system. First, the person is asked to reset the device to rule out false positives. If the person fails to react, a telephone or video-telephone connection will be initiated. Based on location-tracking information, communication devices in the appropriate room will be activated to establish a communication link between the person in trouble and the emergency call centre.

Our approach to reasoning on sensor data raises the challenge of coping with inherently unreliable and imprecise data. This is tackled by combining alternative or redundant sources. For instance, location tracking can be fed from different sensor systems (RFID, ultra-sonic, pressure, etc.).

Furthermore, the system must evolve over time, so as to adapt to the specific conditions and demands of the assisted person, as they acquire new diseases or handicaps or recover from impairment.

The central module of amiCA is MonA (Monitoring and Assistance component). In MonA, sensor data is collected and model-based reasoning is performed. The model base entails models of the assisted person, their capabilities and needs, and of the person's daily routine, including recurring medical treatments, as well as models of the living environment, for example devices and their location and capabilities.

Validation

The impact of our solution will be evaluated in several steps. The first step is the prototypical development and integration in an apartment-like AAL Lab (late 2006). The second step is integration into a nursing home and with the emergency call centre in Kaiserslautern (in 2007). The third step is validation of emergency monitoring within the 6 FP project EMERGE on a pan-European level (2006 – 2008).

This work is part of the joint research project BelAmI of Fraunhofer IESE, the University of Kaiserslautern, the University of Budapest, the University of Szeged, and the Bay Zoltan Research Foundation, Budapest. Recent achievements are described in detail in the paper: Nehmer, J.; Karshmer, A.; Becker, M.; Lamm, R.: "Living Assistance Systems – An Ambient Intelligence Approach", in Proceedings of the 28th International Conference on Software Engineering (ICSE 2006), Shanghai, China, 2006.

Links:

BelAmi: http://www.belami-project.org
IESE: http://www.iese.fraunhofer.de
University of Kaiserslautern research center on AmI website:
http://www.eit.uni-kl.de/AmI/frame.html?en

Please contact:
Frank Bomarius, Martin Becker, Thomas Kleinberger, Fraunhofer IESE, Germany
Tel: +49 631 6800 1200
E-mail: {bomarius, mbecker, kleinberger}@iese.fraunhofer.de

iCat - A Friendly Robot that Helps Children and Grown-Ups

by Bernt Meerbeek, Jettie Hoonhout, Peter Bingley and Albert van Breemen

Can a robot cat be a buddy for children and adults? Philips Research in the Netherlands developed iCat, a prototype of an emotionally intelligent user-interface robot. It can be used as a game buddy for children or as a TV assistant, or play many other roles.

The AmI Paradigm

Ambient Intelligence (AmI) refers to electronic environments that are sensitive, adaptive, and responsive to the presence of people. Its focus is not only the physical integration of electronics, but also the creation and generation of enhanced experiences. Intuitive interaction and emotional user-interface robots will play an important role in the realisation of Ambient Intelligence. Philips Research developed iCat, a research prototype of an emotionally intelligent robot that can provide an easy to use and enjoyable interaction style in AmI environments.

Developed in the late 1990s, the AmI paradigm presents a vision for digital systems from 2010 onwards. Current technological developments enable the integration of electronics into the environment, thus providing people with the possibility to interact with their environment in a seamless, trustworthy and easy to use manner. This implies that embedding-through-miniaturisation is the main systems design objective from a hardware point of view. In software we study context awareness, ubiquitous access, and natural interaction. The user benefits are aimed at improving the quality of peoples' lives by creating a desired atmosphere and providing appropriate functionality by means of intelligent, personalised and interconnected systems and services.

The iCat Concept

Emotional robots are generally considered as a new and promising development for the intuitive interaction between users and AmI environments. iCat was developed as an open interactive robot platform with emotional feedback to investigate social interaction aspects between users and domestic robots with facial expressions; see Figure 1.

iCat is a 38cm tall cat-like robot character. The head contains 13 servomotors that can move the head and different parts of the face, including eyebrows, eyes, eyelids, and mouth. The servomotors can generate facial expressions, which give the robot socially intelligent features; see Figure 2. Through a camera in the nose, iCat can recognise objects and faces using computer vision techniques. Each foot contains a microphone that can identify sounds, recognise speech, and determine the direction of the sound source. A speaker in the bottom can play sounds (WAV and MIDI files) and connected speech. iCat can be connected to a home network to control domestic appliances, and to the Internet to obtain information. iCat can sense touch through sensors in its feet. It can communicate information encoded by coloured light through multi-color LEDs in its feet and ears. For instance, the LEDs in the ears can indicate different modes of operation such as sleeping, awake, busy, and listening.

Psychology meets Technology

One of the key research questions for iCat is to find out whether facial expres-

Figure 1: iCat, here as a game buddy for children.

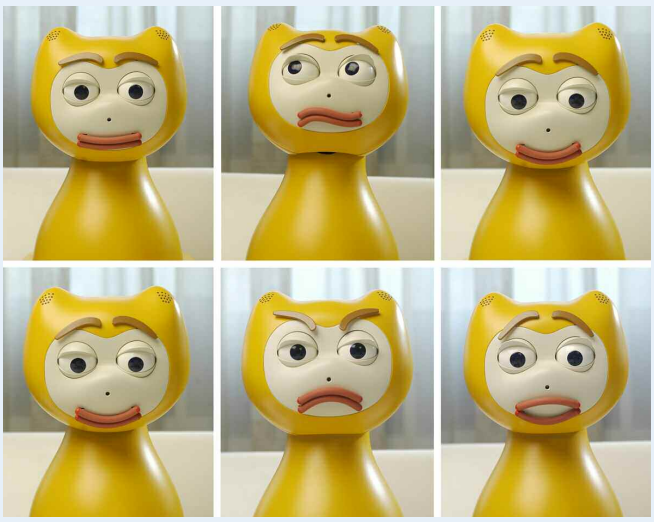

Figure 2: iCat facial expressions.

Pictures by courtesy of Philips.

sions and a range of behaviours can give the robot a personality and which type of personality would be most appropriate during interaction with users for certain applications. Research with iCat has focused on the evaluation of application concepts for iCat, for example as a game buddy for children or as a TV assistant. The research questions in these studies were: "What personality do users prefer?" and "What level of control do they prefer?"

The results of the user study with the iCat as game buddy indicated that children preferred to play games with the iCat rather than to play these same games with a computer. They were able to recognise differences in personalities between differently programmed versions of iCat. Overall, the more extrovert and sociable iCat was preferred to a more neutral personality.

Young and middle-aged adults were also able to recognise differences in personality in iCat as a TV assistant. In this study two personalities were combined with two levels of control. In the high control condition, iCat used a speech-based command-and-control interaction style, whereas in the low control condition it used a speech-based, system-initiated natural language dialogue style. The preferred combination was an extrovert and friendly personality with low user-control.

One of the most interesting results was that the personality of the robot influenced the level of control that people perceived. This is very relevant in the context of intelligent systems that work autonomously to take tedious tasks out of the hands of humans. It suggests that the robot's personality can be used as a means to increase the amount of control that users perceive.

Conclusion

The studies with iCat have shown that mechanically-rendered emotions and behaviours can have a significant effect on the way users perceive – and interact with – robots. Moreover, users prefer to interact with a socially-intelligent robot for a range of applications, compared to more conventional interaction means. A range of further studies is planned. One of the questions that will be addressed is how iCat should behave to inspire trust and compliance in users, important if one thinks of the robot cat as a personal (health) trainer, for example.

Links:
http://www.hitech-projects.com/icat

Please contact:
Bernt Meerbeek, Philips Research Europe, The Netherlands
Tel: +31 40 27 47535
E-mail: bernt.meerbeek@philips.com

Albert van Breemen, Philips Research Europe, The Netherlands
Tel: +31 40 27 47864
E-mail: albert.van.breemen@philips.com

The DECOS Tool-Chain: Model-Based Development of Distributed Embedded Safety-Critical Real-Time Systems

by Wolfgang Herzner, Bernhard Huber, György Csertan and Andras Balogh

'Smart systems' applications like adaptive cruise control or brake-by-wire, rely on predictable and reliable embedded system platforms as infrastructure. For the development of such dependable applications, it is therefore of crucial interest to avoid faults during design and development. Besides intensive testing, an important way of minimising the risk of faults is by controlling the design and development process, as well as maximising the coherence of the resulting system with the initial requirements. The model-based tool-chain developed by the DECOS project is described.

Today, the development of embedded - and in particular safety-critical - systems in general follows a customized design approach, resulting in rather isolated applications and little reuse of components and code across different application domains. For instance, in modern cars sub-systems like power-train control, advanced driver assistance systems or the body electronic co-exist, each equipped with its own electronic hardware, communication cabling etc. This approach implies at least increased hardware costs, weight, and power consumption, last not least due to severely hampering the sharing of resources like sensors among the different sub-systems.

Therefore, the European project DECOS aims at developing basic enabling technology for moving from federated to integrated distributed architectures in order to reduce development, validation and maintenance costs, and increase the dependability of embedded applications in various application domains. 'Integrated' means, that several software 'IP'-blocks (Intellectual Property) of different criticality can be allocated to one node (ECU – Electronic Control Unit) without interfering with each other, ie, guaranteed encapsulation in space (memory) and time (each job has its reserved time slot).

DECOS presumes the existence of a core architecture providing the core services:
- deterministic and timely message transport
- fault tolerant clock synchronisation
- strong fault isolation
- consistent diagnosis of failing nodes.

Any core architecture providing these services (eg TTP/C, FlexRay, or Time-

Triggered Ethernet) can be a basis for DECOS-based systems. On top of these core services, DECOS provides a set of architectural (or high-level) services:

- virtual networks (VN) and gateways
- an encapsulated execution environment (EEE)
- diagnostics.

To minimize the dependency of application programming on a certain DECOS implementation, a Platform Interface layer (PIL) provides a techology invariant interface of the high level services for application tasks.

DECOS Tool-Chain

A constituent element of such an enabling technology is a tool-chain, currently being developed by DECOS, which encompasses all embedded software design and development aspects, including configuration and testing. As illustrated in the figure, the DECOS tool-chain essentially consists of two vertical 'lanes': on the left side, the integrated system configuration is determined and middleware is generated, and on the right side the application functionality is developed. A third lane, containing tools for testing and verifying the various (intermediate) results, is not shown.

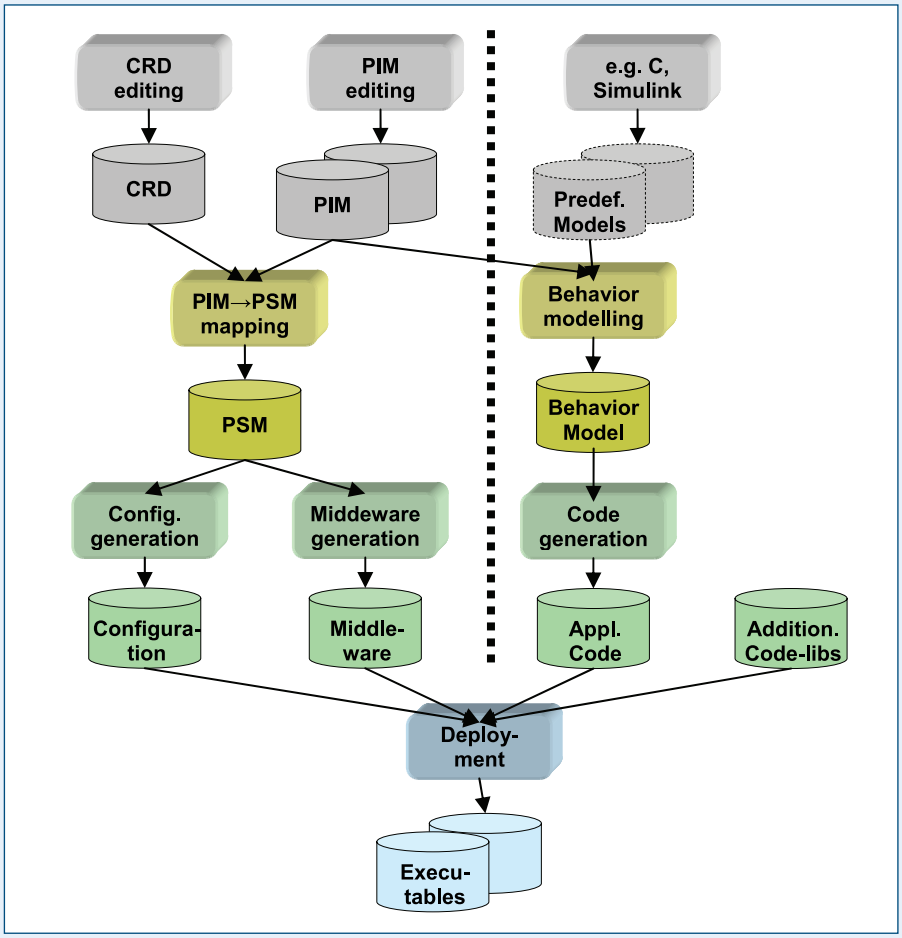

DECOS tool-chain basic steps. Boxes denote activities, and disk symbols denote data. Grey elements address specification, yellow address design, green ones implementation, and the blue elements installation.

The specification starts with the Platform Independent Models (PIMs) of the application sub-systems, defining their requirements with respect to communication (among the application tasks), performance, and dependability.

PIMs serve two purposes: firstly, together with the specification of the target cluster hardware and resources, the Cluster Resource Description (CRD), they are used to derive the Platform Specific Model (PSM), which contains allocation (of tasks to nodes) and other information relevant for the successive steps. From PSM, configuration files and schedules for both task execution and message transmission are generated, as well as middleware like the PIL.

Secondly, PIMs are used to guide the development of jobs (ie application tasks), by modelling their behaviour with SCADE (a tool set of Esterel Technologies). If feasible, predefined Simulink models or modules written in conventional languages like C or Ada can be imported. After application code is generated from these models, the results of both activities are integrated to achieve the target executables, which can then be downloaded to the application cluster.

The purpose of the CRD (Cluster Resource Description) is to capture the characteristics of the platform relevant for the software-hardware integration. This includes computational resources (CPU, memory), communication resources and dependability properties. A graphical domain-specific modeling environment is developed, based on GME (Generic Modeling Environment). The targeted modeling domain is described formally via the HSM (Hardware Specification Model), a meta model which facilitates the validation and reuse of resource models.

Before generating the PSM, it is possible to add information manually to the PIM (PIM marking), for example information on specific middleware requirements.

Jobs are assigned to nodes taking into account functional and non-functional constraints. In the first phase of assignment, constraints are dealt with one by one. Since allocation is an NP-hard problem, in a second phase a multi-variable optimisation approach is proposed.

Scheduling is the next step, where a tool suite (TTplan, TTbuild) of TTTech (a DECOS partner developing time-triggered systems) has been adapted to

handle resource restrictions and EEE partitioning. Then, PIL is generated, providing generic message transfer, global time service and membership service (necessary to distribute information on the state of nodes).

For behaviour modelling, SCADE (by Esterel Technologies) has been chosen as a primary tool for DECOS. SCADE is based on a formally-defined data flow notation. It offers strong typing, explicit initialisation, explicit time-management and simple expression of concurrency. The PIMs are imported via the SCADE UML gateway, yielding empty job skeletons with correct interfaces. Their behaviour is then directly modelled with SCADE, or Simulink models are imported to SCADE via another gateway. SCADE is used for code generation. For linking job code with PIL, so-called 'SCADE-wrappers' are also generated.

The DECOS tool-chain comprises a wide variety of tools from model to deployment. To ease handling, a transformation tool VIATRA, developed at Budapest University of Technology and Economics, is used as backbone for model transformations (from PIM to PSM), PIL generation and domain-specific editors. Four tools are used for the DECOS tool-chain: GME, VIATRA, SCADE and TTplan/build; additionally, commercial and target-platform specific tools are used for deployment (compilation, linking, download). This tool-chain is designed for efficient configuration, development and validation of critical 'smart' embedded applications.

Links:
http://www.smart-systems.at
http://www.decos.at

Please contact:
Wolfgang Herzner, ARC Seibersdorf research/AARIT, Austria
Tel: +43 50550 4132
E-mail: wolfgang.herzner@arcs.ac.at

Bernhard Huber, Vienna University of Technology, Austria,
E-mail: huberb@vmars.tuwien.ac.at

György Csertan, Andras Balogh, Budapest University of Technology and Economics
E-mail: {csertan, abalogh}@mit.bme.hu

Ambient Hardware: Embedded Architectures on Demand

by Kjetil Svarstad

Imagine you are at the airport, your flight is indefinitely delayed, you are bored, and your PDA-based chess program offers you no real challenge. Luckily, you have a machine description of the most powerful chess machine architecture (like IBM's Deep Blue) available on your PDA which you download to the nearest AHEAD-tag server, and then you can play Grand Master-level chess on the tag with your PDA as client until your flight is finally called.

Ambient Intelligence (AmI) is considered the next huge leap in computer and information technology. One research project in this area, 'Ambiesense', used small computational tags serving mobile clients with information and computation via wireless services. Our project, AHEAD, integrates the concept of reconfigurable electronics with the computational tags, enabling mobile clients to hold actual configuration data and reconfigure the tags on demand through a wireless service.

The main idea of AHEAD - Ambient Hardware: Embedded Architectures on Demand - is the combination of three technological evolutions:

- Hardware or computational machines becoming smaller and disappearing into the surroundings.
- Ambient Intelligence, the disappearing computers being able to communicate and in concert solve problems in new ways.
- Reconfigurable hardware, FPGA (Field-Programmable Gate Array, a type of logic chip that can be programmed), closing the gap between hardware and software, making functionality truly mobile between machine and program.

In the AmI perspective computers will be tightly embedded in the surrounding environment and the computational tasks and services adaptable to any and many different situations and thus architectures. Our project hypothesis extends this view of embedding the hardware and adapting the software into adaptable and movable hardware, what we have chosen to call 'Ambient Hardware'. This is a platform where small computers (tags) in the environment are able to reconfigure themselves in order to fulfill the combined hardware-software requirements of a given task. The underlying architecture is tag computers equipped with reconfigurable coprocessors in terms of FPGAs. Physically the tag computer is also implemented on the FPGA, and it reconfigures itself as the requirements change as a result of clients and their tasks entering and leaving the activity zone of the tag.

Although research projects and results abound in the areas of Ambient Intelligence and Reconfigurable Computing, our objectives in AHEAD were best solved through practical experiments to show which of our ideas were sound, subsequently forming the underlying theories when we know specifically which architectures and solutions are the most promising. Presently we are still in the experimental stage of the project.

Our first objective was to choose a good FPGA platform for our laboratory exper-

iments. We tested and evaluated several development boards for different FPGAs. However, the flexibility of reconfiguration was a limitation on most of them. In the end we chose a very simple and available FPGA board, the so-called 'Suzaku system', which was small and allowed us the flexibility of reconfiguration that we required. This system is really a Single Board Computer in an FPGA with an external flash memory for reconfiguration data and some communication controllers for LAN and serial line I/O.

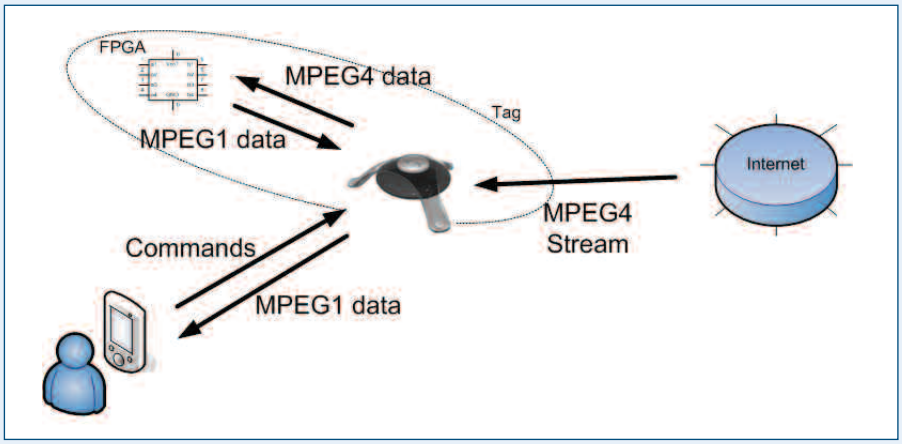

Media streaming execution.

We are now at version 1.0 of our project roadmap. This is the first major version where the tags are independent of any development PCs or external computers. The tag is now completely embedded on the Suzaku board where the Xilinx Spartan FPGA is configured with a Microblaze CPU and devices for running a Linux operating system variant. Additionally, the tags are independent in the sense that they can reconfigure themselves. A simple web server is running on the tag and, through a specific interface, clients may download new configurations for the FPGA, then the tag will reconfigure itself and come alive again with the basic CPU + Linux configuration plus the newly configured hardware architecture and the interconnections between these sub-systems. The web server interface enables a client to download new configurations either directly from itself or from any networked server holding properly defined system configurations.

In order to test version 1.0 of the AHEAD tag, we are currently developing a demonstrator for MPEG recoding and re-scaling. The objective is to configure the tag for receiving massively encoded media streams such as the MPEG4 format and to decode the stream, downscale it for a small screen low-resolution viewer such as a PDA or a Smartphone and recode the stream in a more easily decoded MPEG format (see illustration). MPEG4-streams are of course well suited to limited bandwidth distribution over the net, but the complexity of decoding them is too demanding for low performance and limited power devices such as PDAs and Smartphones. The tag recodes the stream for less performance and power in decoding, but the format will demand a higher bandwidth. This demand, however, is not a problem since the tag and the client are very close and higher bandwidths are possible even with low power. We expect to test this tag-based recoder version of AHEAD by the end of this year.

Please contact:
Kjetil Svarstad, NTNU, Norway
E-mail: kjetil.svarstad@iet.ntnu.no

System-Level Design of Fault-Tolerant Embedded Systems

by Alain Girault

Fault-tolerance is the ability of a system to maintain its functionality, even in the presence of faults. With the advent of ubiquitous computing and distributed embedded systems, it is becoming an aspect more and more crucial. We have provided new functionalities to the SynDEx system-level CAD software. SynDEx is ideal for optimising distributed real-time embedded systems and our new functionalities allow us to guarantee a specified fault-tolerance level for the generated embeddable code.

Our contribution to research in the fault-tolerant embedded systems consists of several scheduling/distribution heuristics. Their common feature is to take as an input two graphs: a data-flow graph ALG describing the algorithm of the application and a graph ARC describing the target distributed architecture (see figure).

Also shown is a table giving the worst-case execution time of each operation onto each processor and the worst-case transmission time of each data-dependence onto each communication link. The architecture being a priori heterogeneous, these need not be identical. Below is an example of such a table for the operations of ALG. The infinity sign expresses the fact that the operation I cannot be executed by the processor P3, for instance, to account for the requirement of certain dedicated hardware.

From these three inputs, the heuristic distributes the operations of ALG onto the processors of ARC and schedules them statically, together with the communications induced by these scheduling decisions. The output of the heuristic is therefore a static schedule from which embeddable code can be generated.

Our fault hypothesis is that the hardware components are fail silent, meaning that a component is either healthy and works fine, or is faulty and produces no output at all. Recent studies on modern hardware architectures have shown that a fail-silent behaviour can be achieved at a reasonable cost, so our fault hypothesis is reasonable.

Our contribution consists of the definition of several new scheduling/distribution heuristics in order to generate static schedules that are also tolerant of a fixed number of hardware components (processors and/or communication links) faults. They are implemented inside SynDEx, as an alternative to its own default heuristics (called DSH: Distribution Scheduling Heuristic):

- FTBAR (Fault-Tolerant Based Active Replication) generates a static schedule that tolerates Npf processor faults by replicating actively all the operations of the algorithm graph ALG exactly Npf+1 times. It works with target architectures having either point-to-point communication links or buses, but assumes that all the communication links are reliable. FTBAR tries to minimise the critical path of the obtained schedule w.r.t. the known WCETs of the operations onto the various processors of the architecture.
- RBSA (Reliable Bicriteria Scheduling Algorithm) also generates a reliable and static schedule by actively replicating the operations of the algorithm graph. The difference with FTBAR is that the number of times an operation is replicated depends on the individual reliability of the processors it is scheduled on and on the overall reliability level required by the user. RBSA tries both to minimise the critical path of the obtained schedule and to maximise its reliability (these are the two criteria of this heuristic).

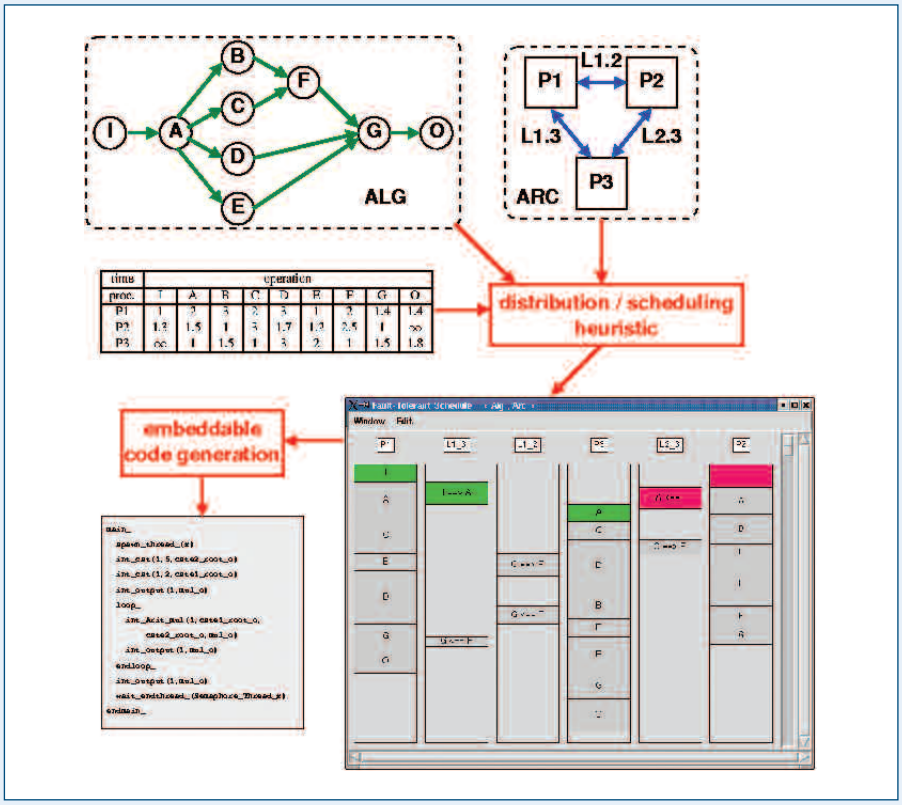

To the left is an example of an algorithm graph: it has nine operations (represented by circles) and 11 data-dependences (represented by green arrows). Among the operations, one is a sensor operation (I), one is an actuator operation (O), while the seven others are computations (A to G). Below to the right is an example of an architecture graph: it has three processors (P1, P2, and P3) and three point-to-point communication links (L1.2, L1.3, and L2.3).

- GRT + eDSH (Graph Redundancy Transformation + extended Distribution Scheduling Heuristic) generates a static schedule that tolerates Npf processor faults and Nlf communication link faults. It first transforms the algorithm graph ALG into another dataflow graph ALG* by adding redundancy into it such that the required number of hardware component faults will be tolerated. During this phase, it also generates exclusion relations between subsets of operations that must be scheduled onto distinct processors, and subsets of data dependences that must be routed through disjoint paths. Then it uses an extended version of the DSH heuristics to generate a static schedule of ALG* onto ARC, w.r.t. the exclusion relations generated during the first phase.
- FPMH (Fault Patterns Merging Heuristic) is an original approach to generate a static schedule of ALG onto ARC, tolerant to a given list of fault patterns. A fault pattern is a subset of the architecture's component that can fail simultaneously. Our method involves two steps. First, for each fault pattern, we generate the corresponding reduced architecture (the architecture from which the pattern's component has been removed) and we generate a static schedule of ALG onto this reduced architecture (we use the basic DSH heuristic of SYnDEx for this). From N fault patterns we therefore obtain N basic schedules. The second step consists of the merging of these N basic schedules into one static schedule that will be, by construction, tolerant to all the specified fault patterns.

Links:

Fault-tolerance:
http://pop-art.inrialpes.fr/~girault/Projets/FT/
SynDEx: http://www.syndex.org

Please contact:
Alain Girault, INRIA Rhône-Alpes
Tel: +33 476 61 53 51
E-mail: Alain.Girault@inrialpes.fr

New Tool to Design the Behaviour of Embedded Systems

by Jan Friso Groote

Can a car drive faster than the speed limit the cruise control indicates? Are computerised protocols between planes and air traffic control safe? Communication between software components in cars, planes and other intelligent embedded systems is often very complex and prone to errors. The Design and Analysis of Systems group at Eindhoven University of Technology is releasing a new toolset, mCRL2, to improve the quality of embedded system design.

An embedded system can roughly be described as some piece of equipment with one or more computer processors in it. The processor allows this equipment to behave with an amazing level of intelligence. Copiers order their own paper and cars actively avoid accidents. The end of this is not in sight. The behaviour of such systems will become far more advanced than it is today.

The downside of this is that the behaviour of systems becomes so complex that we cannot easily understand it anymore. As a consequence, most features of contemporary equipment remain unused. But it is even worse. Because the system components are exhibiting more and more complex behaviour, they do not always understand each other. The consequence is that intelligent systems often behave erraticly. For modern TV sets a 'misunderstanding' between the components occurs every minute in each television. Much effort goes into building internal firewalls to restrain these flaws so they are not noticed by those watching TV.

Now that the behaviour of systems is becoming so complex, it has become an important research topic. Just as in other engineering disciplines, complexity is tackled by making models. These behavioural models describe the potential behaviour of the system, ie which interactions can take place when. Typical examples are sending messages, pressing buttons, activating actuators and reading sensors. Questions that must be answered are: are all messages properly processed; can a car ever be instructed to drive faster than the speed limit of the cruise control; does a copier always order paper in time or is it possible that it orders twice the amount needed due to message duplication?

The major challenge to answering such questions is the 'state space explosion'. Even relatively simple behavioural models give rise to millions of states. In order to show that a system behaves correctly, all these states must be investigated. It goes without saying that analysis tools are essential here. It is impossible to carry out the analysis of an industrial behavioural model by hand, but manual manipulation and human insight is also essential to reduce it to a manageable size.

At Eindhoven University of Technology, in close cooperation with CWI in Amsterdam, we are developing a mathematical methodology accompanied with tools to model and analyse the (discrete) behaviour of embedded systems. The method is called mCRL2, where CRL stands for Common Representation Language. The methodology is based on process algebra, extended with data types. The method has been applied to almost all major embedded system industries within the Netherlands. Furthermore, it is being taught at several universities.

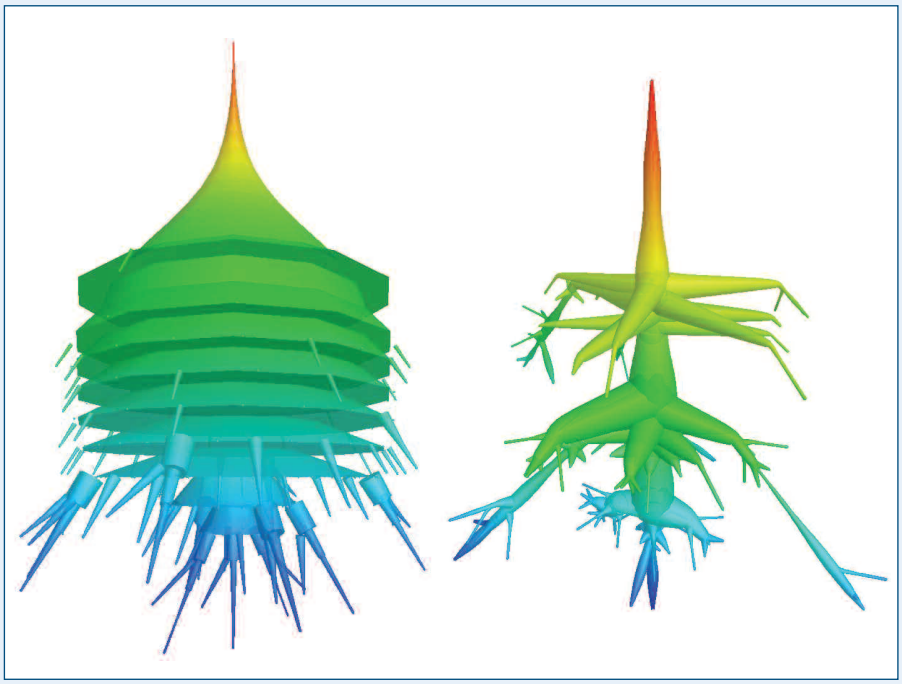

Picture: TU/e.

State spaces, visualized by the mCRL2 toolset.

The major idea underlying the methodology is to transform each model to a normal form, a so called linear process. Due to its simple structure, a linear process is very amenable to symbolic manipulation and simplification. For instance, by detecting confluent

behaviour in a linear process, the associated state space is often exponentially smaller than without the use of confluence. Experience has shown that all models can straightforwardly be translated to linear form. Unfortunately, not every form of analysis can be carried out on linear processes. For this state automata are needed, but these can not always be generated.

Besides the analysis tools, the toolset comes with several visualisation tools. In particular, tools that can visualise the structure of state spaces of up to a million states are very helpful in obtaining an insight into the behaviour of embedded systems. These techniques have been used more than once to detect symmetric or erroneous behaviour. They have also been very helpful in understanding the extent of problematic behaviour compared with the overall behaviour.

The tools and techniques are being developed in the scientific domain. This not only means that the ideas and technologies being used are published within the scientific community (as is common), but also that the tools are an open source, free for both scientific and commercial use. A preliminary version of the toolset can be found at www.mcrl2.org. A full first release is expected in the autumn of 2006 and will be available for Linux, Mac OS X and Windows. It will contain all available modelling, analysis and visualisation tools.

Please contact:
Jan Friso Groote, Eindhoven University of Technology, The Netherlands
Tel: +31 40 2475003
E-mail: jfg@win.tue.nl
http://www.win.tue.nl/~jfg

Reactive Processing for Reactive Systems

by Reinhard von Hanxleden, Xin Li, Partha Roop, Zoran Salcic and Li Hsien Yoong

Embedded real-time systems must react continuously to stimuli from their environment. Therefore, their control-flow patterns differ from those of traditional systems that transform a given input to an output at their own pace. Reactive processors provide direct hardware support for reactive control flow, which keeps executables fast and compact and results in lower power consumption compared to traditional architectures.

Many embedded systems belong to the class of reactive systems, which continuously react to inputs from the environment by generating corresponding outputs. The programming of reactive systems typically requires the use of non-standard control flow constructs, such as concurrency and exception handling. Most programming languages, including languages such as C and Java that are commonly used in embedded systems, either do not support these constructs at all, or their use induces non-deterministic program behaviour, regarding both functionality and timing. Furthermore, the concurrent behaviour of reactive programs is typically very fine-grained. Measurements indicate that it is not uncommon to have a context switch after, on average, fewer than five instructions. This makes traditional context switch mechanisms, based on an operating system that consumes thousands of instructions per context switch, impractical.

To address these difficulties, the synchronous language Esterel has been developed to express reactive control flow patterns in a concise manner, with a clear semantics that imposes deterministic program behaviour under all circumstances. There are different options to synthesise an Esterel program into a concrete system. In the software synthesis approach, the Esterel program is translated by an Esterel compiler into a C program, which in turn runs on a COTS processor. Esterel can also be synthesised directly into hardware, via some hardware description language such as VHDL. As a hybrid approach, Esterel can also be used in HW/SW co-design. However, these classical synthesis approaches suffer from the limitations of traditional processors, with their instruction-set architectures geared towards the sequential von-Neumann execution model, or they are very inflexible if HW synthesis is involved.

Recently, another alternative for synthesising Esterel has emerged, the reactive processing approach. Here, the Esterel program is running on a processor that has been developed specifically for reactive systems. The instruction set of these reactive processors closely resembles the constructs found in Esterel, such as waiting for the occurrence of a signal or abortion. A comparison highlighting the main philosophical differences between reactive processors and conventional processors is outlined in Table 1.

The intention of reactive processors is to unite the pros of both the software and hardware synthesis approach while avoiding most of the cons of both. We get from the software synthesis approach the easy programming process with inexpensive testing and short turnaround times. From the hardware side, we gain low energy consumption and fast program execution. The synchronous model of computation inherent in Esterel ensures deterministic behaviour.

Two strategies have been proposed to design reactive processors, which mainly differ in the way they handle concurrency. Since Esterel is a synchronous language, the operation of all concurrent

threads is based on the ticks of a single global clock. The principal challenge is to schedule concurrent threads so that within a logical tick, event producers are executed before event consumers. The EMPEROR processor, developed at the University of Auckland, New Zealand, employs multi-processing and uses dynamic scheduling of events. This allows it to handle many threads running on a cyclic executive on a given processor or between concurrent threads running on different processors. The Kiel Esterel Processor (KEP) family, developed at Kiel University, Germany, handles concurrency via multi-threading, which minimises overall resource usage and easily scales up to very high degrees of concurrency. Unlike the dynamic self-scheduling of threads in the EMPEROR approach, the KEP approach schedules all producers before consumer threads to uniquely determine signal statuses. Both approaches have their advantages and limitations. However, overall they perform much better than conventional processors while performing reactive computations.

Apart from efficiency and determinism concerns, another advantage of reactive processors is that due to their comparatively simple structure (no caches, no pipelining) and their direct implementation of reactive control flow constructs, it becomes feasible to precisely characterise their timing behaviour. In conjunction with the synchronous model of computation, which discretises time into logical ticks, it is possible to derive exact, tight bounds on its Worst Case Reaction Time (WCRT), which tells how much time it takes the system to react to the environment. The KEP processor series is equipped with a Tick Manager that can provide a constant logical tick rate and also detects internal timing over-runs. This can serve to detect hardware failures and provides another safeguard, in addition to static analyses, so that real-time deadlines are met.

Features	Reactive Processors	Conventional Processors
Execution progression	Evolves in discrete instants separated by "tick delimiting instructions".	Evolves continuously.
Preemption	Accomplished through event reaction block with implicit priority resolution and context switching in hardware.	Accomplished through interrupt mechanism requiring explicit priority resolution, context saving and restoration in software.
View of the environment	Changes at discrete instants. Inputs are latched at the beginning and outputs are sustained till the end of a "tick".	Changes continuously. Inputs can be read at any time, and outputs can be sustained for any duration.
Concurrency	Synchronous parallel execution and broadcast communication between threads.	Asynchronous execution requiring explicit message passing/rendezvous for communication between threads.

Table 1: Comparison between reactive and conventional processors.

FPGA-based implementations of reactive processors have proved very competitive to classical processor designs. For a standard suite of Esterel benchmarks the code size is typically an order of magnitude smaller than that of the MicroBlaze, a 32-bit COTS RISC processor core. The worst case reaction time is typically improved four-fold and energy consumption is also typically reduced to a quarter.

In the five years since its inception, the reactive processing approach has thus demonstrated its promise and its practicality. However, much remains to be done. On the theoretical side, a precise characterisation of the reactive execution semantics is still missing and its relationship to other semantics needs to be investigated, in particular regarding causality issues. The reactive instruction set architectures also pose interesting compiler challenges, for example, regarding the efficient mapping of concurrency with minimal context switching. Finally, the derivation of WCRT bounds is so far done very conservatively - it should be feasible to tighten these analyses further.

Links:

http://www.informatik.uni-kiel.de/rtsys/kep

http://www.ece.auckland.ac.nz/~roop/pageGen.php?pageid=10

Please contact:

Reinhard von Hanxleden, Christian-Albrechts-Universität zu Kiel, Germany

Tel: +49 431 880 7281

E-mail: rvh@informatik.uni-kiel.de

Partha Roop, University of Auckland, New Zealand

Tel: +64 9 3737599

E-mail: p.roop@auckland.ac.nz

TPC: Tagged Procedure Calls

by Konstantinos Kapelonis, Sven Karsson, and Angelos Bilas

Tagged Procedure Calls (TPC) is a new approach addressing the problem of the high programmer effort needed to achieve scalable execution. TPC is targeted at architectures ranging from small embedded systems to large-scale multi-core processors and provides an efficient programming model easy to understand and exploit.

A huge opportunity and challenge we face today is the design of embedded systems that will support demanding application domains. Current technology trends in building such embedded systems advocate the use of parallel systems with (i) multi-core processors and (ii) tightly-coupled interconnects. However, exploiting parallelism has traditionally resulted in significant programmer effort. The programming model plays an important role in reducing this effort. The main challenges in the programming model are to expose mechanisms that need to be used directly by the programmer and to hide mechanisms that can be used transparently.

It is predicted that multi-core processors will use the increasing number of transistors more efficiently than traditional single-core processors and offer higher performance. However, multi-core processors require a large degree of parallelism that has traditionally demanded additional programming effort. There is a need for an intuitive programming model for multi-core processors.

Moreover, future interconnects will play an important role in such systems as they interconnect all the cores in a system. To be efficient, the programming model must facilitate efficient use of the interconnect.

Tagged procedure calls (TPC), are a new programming model that re-draws the balance between what the programmer needs to specify and what the architecture should provide. Furthermore, TPC aim at unifying intra- and inter-chip programming abstractions. Informally, the main points behind TPC are:

- TPC allow the programmer to augment procedure declarations and calls with tags that control the way these procedures will be executed.
- TPC require the programmer to specify parallelism through special, tagged, asynchronous procedure calls. TPC also provide primitives that block execution until certain procedure calls have completed, see Figure 1.
- TPC discourage the programmer from accessing global data within tagged procedures.
- TPC allow the programmer to use pointers to global data as arguments to procedures.

```
/* Function with OUT argument */
void procedure(...){
...
}
void main()
{
 /* Define a handle for function */
 TPC:DECL(hdl,procedure)
 ...
 /* call asynchronously procedure()
  * it changes argument s */
 TPC:CALL(hdl,procedure,s) c,i
 /* Continue execution */
 ...
 /* Use handle hdl to wait for
  * procedure() to complete */
 TPC:WAIT(hdl)
 /* s can now be accessed */
 ...
}
```

Figure 1: TPC example. Calling a function with consistency (c) and isolation (i) tags.

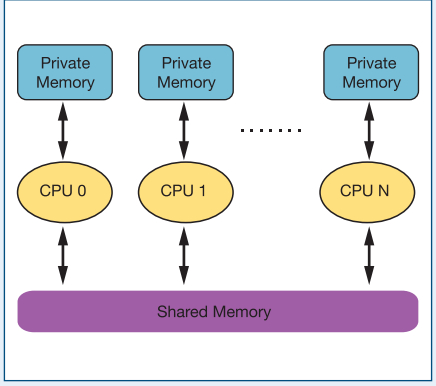

Figure 2. Current prototyping platform at FORTH-ICS. The depicted architecture is implemented using a Xilinx XUP board.

Tags placed on procedure calls by the programmer define an abstract representation of the required execution semantics, see Figure 1. The programmer (a) has to identify the available parallelism, (b) is urged to identify the data used during parallel computations and (c) has to specify how parts of the code will execute, eg as atomic or serializable regions. However, they need not worry about how this will be achieved. Finally, TPC aim at unifying intra- and inter-processor programming models and dealing with the inherent heterogeneity of future embedded systems in a simple and intuitive manner.

TPC encourage the programmer to specify parallelism and data used during parallel execution, but do not expose the underlying mechanisms for communication, synchronization, etc. We believe that this balance between what the programmer needs to specify and what can be done transparently will result in both efficient execution and reduced programmer effort on future embedded systems.

Our current work focuses on a prototype implementation of TPC on an embedded platform. We are using an FPGA-prototype (see Figure 2) with multiple cores. Using this implementation platform, we are currently experimenting with the semantics of TPC. We are porting a number of existing parallel applications. This effort will allow us to examine the cost introduced by the runtime system when transparently dealing with communication and synchronization issues.

This work has been partially supported by the European Commission in the context of the SARC project.

Please contact:
Angelos Bilas, FORTH-ICS and University of Crete, Greece
E-mail: bilas@ics.forth.gr
http://www.ics.forth.gr/~bilas/

A European Ambient Intelligence Research Facility at FORTH - ICS

by Constantine Stephanidis

Ambient Intelligence (AmI) presents a vision of tomorrow where 'intelligent' environments react in an attentive, adaptive and active (sometimes proactive) way to the presence and activities of humans and objects in order to provide appropriate services to the inhabitants of these environments. It is an emerging field of research and development that is rapidly gaining wide attention from an increasing number of researchers and practitioners worldwide, particularly in Europe.

Ambient Intelligence technologies integrate sensing capabilities, processing power, reasoning mechanisms, networking facilities, applications and services, digital content and actuating capabilities distributed in the surrounding environment. While a wide variety of different technologies is involved, the goal of Ambient Intelligence is to either entirely hide their presence from users or to smoothly integrate them in their surroundings as enhanced environment artifacts rather than technological gadgets. This way, the computing-oriented connotation of technology essentially fades out or even disappears in the environment, providing seamless and unobtrusive interaction paradigms. Therefore, people and their social situation, ranging from individuals to groups, be they work groups, families or friends and their corresponding environments (office buildings, homes, public spaces, etc) are at the centre of the design considerations.

Ambient Intelligence brings a special perspective to the on-going research associated with technical fields like ubiquitous computing, pervasive and proactive computing, ambient computing, embedded computing, and smart objects. Ambient Intelligence has become well-focused by putting people and social contexts at the centre, while at the same time aiming to distribute, embed, coordinate and interactively deliver computing intelligence within the surrounding environment. The notion of Ambient Intelligence, as described above, is becoming a de facto key dimension of the emerging Information Society, since many of the new generation industrial digital products and services are clearly shifted towards an overall intelligent computing environment.

From a technological point of view, Ambient Intelligence is a particularly complex, multi-faceted and demanding scientific domain, requiring the presence and seamless integration of most of the key technologies existing today, thus posing numerous challenges in several research areas and requiring large scale experimentation.

The Institute of Computer Science (ICS) of FORTH is in the process of creating a large-scale, state-of-the-art Ambient Intelligence European Facility, which will act as a research nexus for studying and developing, from a human-centred perspective, related technologies and for assessing their impact on the individual, as well as on society as a whole.

Starting with the creation of an intelligent home simulator, the AmI facility will be expanded to address specific indoor and outdoor environments, relevant application domains (eg housing, work, health, security, education, transportation and entertainment), as well as their related physical, social and cultural characteristics.

The AmI Facility will also provide a showcase for demonstrating the potential added-value and benefits of AmI technologies in different aspects of everyday life and activities. In this way, the AmI Facility will foster the vision of Ambient Intelligence, facilitate multi-disciplinary international collaborations and provide a focal point for technology transfer and dissemination of know-how to European industry, adding to its competitive advantage. It will also contribute towards the European strategic priority for "an Information Society that is inclusive, provides high-quality public services and promotes quality of life", while also promoting synergies and knowledge diffusion in the context of relevant European research projects and actions.

The creation of this facility builds on the scientific know-how and technical expertise of FORTH-ICS in a number of contributing critical domains such as Human-Computer Interaction, Universal Access, Artificial Intelligence, Semantic-based Knowledge Systems, Robotics, Computational Vision, Networks and Telecommunications, Information Security, Distributed Systems, Computer architecture, Microelectronics, Sensors and Biomedical Informatics. All Laboratories of FORTH-ICS are actively engaged in this effort, and collaborations are promoted with other FORTH Institutes, as well as other organisations within and beyond ERCIM.

In the context of promoting research in the domain of AmI, FORTH-ICS has played a key role in the establishment of the new ERCIM working group: 'Smart Environments and Systems for Ambient Intelligence'. The AmI Facility of FORTH-ICS will promote and support active collaboration and synergies among the members of this ERCIM working group, by offering a technological platform and an experimental testbed for research and development activities.

Links:

SESAMI Working Group:
http://www.ics.forth.gr/sesami/
ICS-FORTH Ambient Intelligence Programme: http://www.ics.forth.gr/ami

Please contact:

Constantine Stephanidis, FORTH-ICS, Greece
Tel: +30 2810 391741
E-mail: cs@ics.forth.gr

Ecological Approach to Smart Environments

by Veikko Ikonen

The project EASE (Ecological Approach to Smart Environments) aims to offer more profound analysis and research-based evidence to support the design of smart environments. Furthermore, the aim is to analyse the design issues related to the smart environments and offer some general guidelines for the design of smart environments in different application areas.

The technological shift of computing, including applications and services directed to two dimensions, embeddedness and mobility, has already changed a lot about our relationship to our environment (both social and technological).

Technology has always been embedded in our living environment in some way. The technological infrastructure has been fading from our sight as technology has reached a more mature status. In the electronic and computer era the wires, base stations and servers are usually hidden and we see only our personal technical appliances. However, the embeddedness of technology is now shifting from particular, computer-situated spaces, towards a computers-everywhere philosophy. The idea that we have a place where our interaction with our environment is supported by computers is changing. Soon we will have technological components (communicating with each other) everywhere, out of our sight, obtaining and utilising information gathered from the environment.

The increasing mobility of information and communication technologies has also changed our relationship with our environment. With mobile computers we can carry lots of data with us (eg books, music and photos), but at the same time we can also create new expressions of ourselves and share this information with others. This connection to the global network enables continuous information-sharing and communication in various ways. When these two dimensions work together, when the person with mobile technology interacts with the situated smart environment, we are approaching the area of ubiquitous (or pervasive) computing, also called ambient intelligence.

Even though research in the area of technologically-embedded intelligent environments is expanding very fast, the development of design approaches, other than technology-driven ones, is still in its early stage. The methodologies commonly used in designing and evaluating information and communication technologies have been quite goal-oriented, but new approaches and methods have been introduced and tested along the way. Design principles for intelligent environments have been published, but in these statements the technological issues are also emphasised to a greater extent. More profound co-operation between different designers, developers and researchers is needed in order to put forward new theories and methodologies that will help to develop solutions that naturally support people in their living environment and take into account both the complexity of the systems and technological aspects as well as social, ethical and cultural issues.

With EASE, we have brought together participants from different fields in designing computerised smart environments and especially those who deal with human technology interaction in this research area. The aim of the project is to identify, model and evaluate the ecology of an intelligent environment from the perspectives of human, technological, environmental and different design methodologies. We are specifying the attributes and characteristics of good smart environments as well as identifying and developing suitable methods for the design, implementation and introduction of technologically-enhanced smart environments. The participants make up an interdisciplinary team, which is vital in order to attain truly holistic knowledge of the phenomena. The project pools the knowledge that research partners have accumulated in other projects. Our aim is to cultivate this know-how and to develop and present it in a more general manner for designing future smart environments.

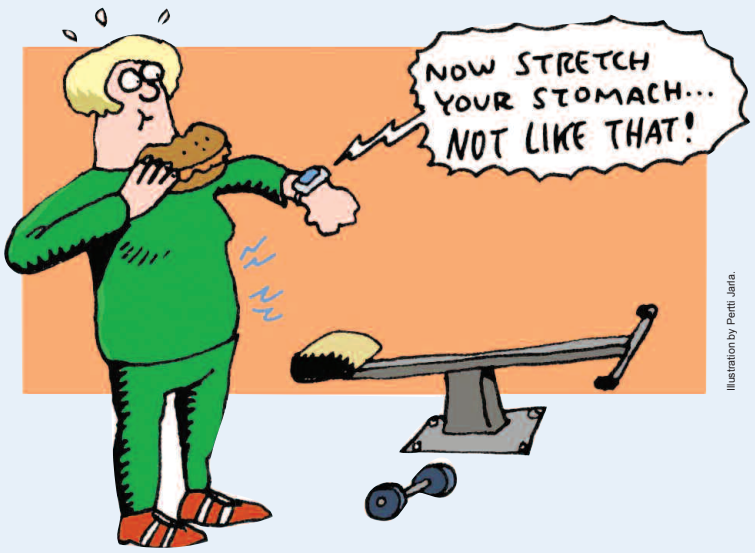

Training with smart technologies.

Ecology is defined by EASE as research of the dynamic interaction system between environment (eg technology in general) and people, in any relation to this technologically augmented environment. Users' actions with the smart applications shape the environment, and the smart environments produce effects on the users and the usage which may even be reflected outside the actual usage environment. The practise of this dynamic relationship is in continuous transition although it seeks stability.

Future computerised smart environments are a challenging design target. This issue is especially tricky when designing public places and multi-user environments. Private or semi-private spaces (ie, work, car, home) can be adjusted more easily according to individual users or a certain user group. It is also easier to compose common rules and regulations for workplaces than for public spaces, for example. One of the great challenges and opportunities would be to integrate the designer and the user again and give control over his computerised environment back to the user.

The technical development and consideration of technological issues of smart environments is going to be a huge job for our global society. Equally important is to consider cultural (including social and ethical) issues related to the technological progress of future computerised environments. User involvement and contextual studies of human technology-interaction in general are going to be a largely accepted curriculum of society in the near future. The great challenge is to maintain a holistic or ecological approach as a part of this curriculum, due to its multi-disciplinary nature.

We want to test our approach in multi-faceted environments where contexts are overlaid, interruptions in simultaneous tasks are more regulation than exception and where various user groups, with different skills and technologies, want to accomplish their primary and secondary goals. Usually, the only valid way to evaluate the concept is in the real-life longitude studies and follow-up studies.

The timeframe of the project is January 2005 to December 2006 and it is funded through the Finnish Funding Agency for Technology and Innovation (Tekes), VTT (Technical Research Centre of Finland) and TTL (Finnish Institute of Occupational Health). Other project partners are the University of Lapland, Tampere University of Technology, University of Tampere, University of Oulu, Helsinki University of Technology and University of Art and Design Helsinki. The project is coordinated by VTT. The steering group of the project consists of a wide range of enterprises and organizations.

Link:
http://www.vtt.fi/proj/ease/

Please contact:
Veikko Ikonen
VTT - Technical Research Center of Finland
Tel: +358 20 722 3351
E-mail: Veikko.Ikonen@vtt.fi

Ambiance: A Platform for Macroprogramming Personalised Ambient Services

by Reza Razavi

A key enabling technology for Ambient Intelligence (AmI) is Wireless Sensor Networks (WSNs). Macroprogramming WSNs by non-programmer end-users is being studied as a step towards an omnipresent World Wide Web interface for the provision of personalised ambient services.

A key enabling technology for AmI is networks of large numbers of wirelessly-connected small, low-powered computers. Such a system is called a Wireless Sensor Network (WSN) and each node, a mote. WSNs can serve as an infrastructure for the provision of personalised ambient services. However, WSNs face very limited processing, memory, sensing, actuation and communication ability of their motes. Programming such networks means those limitations need to be addressed.

Unfortunately, current methods for WSN programming have led developers to mix serious concerns, such as quality of service requirements, with low-level concerns like resource management, synchronisation and routing. This makes developing software for WSNs a costly and error-prone endeavour, even for expert programmers.

Macroprogramming is a new technique which is being developed with the aim of allowing programmers to capture the operation of the sensor network as a whole.

In this research, we focus on simplifying sensor network programming by developing a platform which supports macroprogramming by non-professional programmers. Our aim is to minimise the required programming knowledge, empowering ordinary users to interact with the network so that they can intuitively formulate the expected services. Our architecture also supports an open, concurrent system – requests may come in asynchronously from uncoordinated end-users. They are formulated using an intuitive and omnipresent World Wide Web interface. They are served ubiquitously and in parallel.

Architecture of the Ambiance Platform

To achieve the above goals, we are working on extending the Adaptive Object-Model (AOM) architectural style with Actor-based concurrent computa-

tion, communication, sensing and actuating. The resulting platform, called Ambiance, provides a meta-level architecture to automatically translate high-level specifications of global behaviour by end-users into both meta-objects and meta-actors. These control and customise the runtime behaviour of passive and active application objects.

The meta-objects are dynamic. They have the capability to observe the application objects and the environment (introspection) and to customise their own behaviour by analysing these observations (intercession). Behaviour models and requested services are specified by end-users, in their own terms, using a sophisticated interface provided by Ambiance. Our architecture also allows meta-objects to modify their behaviour in more fundamental ways if, for example, the meta-objects are endowed with learning mechanisms.

The architecture of the Ambiance platform comprises four subsystems (see figure, from left to right):

1. *Presentation subsystem (APE):* the system's front-end. It is a web-enabled application which is responsible for the presentation and interaction with end-users.
2. *Representation subsystem (ARE):* offers to ordinary end-users an intuitive, spreadsheet-like task-specific language for commanding the system. Web pages sent to the APE are dynamically generated by ARE, in HTML, thanks to the Seaside framework and Squeak. ARE exports service models in XML to the Transformation subsystem (see below). It also sends execution requests to the Deployment subsystem and receives from it the execution result. On reception, the latter is communicated to the APE (for formatting and presentation to the end-users).
3. *Transformation subsystem (ATE):* implemented in VisualWorks Smalltalk, is responsible for model transformation and verification. It allows service requests to be deployed on different runtime environments.
4. *Deployment subsystem (ADE):* is responsible for executing models in an adaptive, distributed, and concurrent environment.

For implementing this architecture we reuse the Dart meta-level object-oriented framework for task-specific, artifact and activity-driven behaviour modeling. This framework offers the reifications needed for (1) explicitly representing the ambient services and (2) automating the selection and deployment of an appropriate execution strategy, according to the service's resource consumptions and the actual execution environment (Context-aware Computing). Each specification is translated into a group of meta-actors which implement protocols to meet it.

These meta-actors control the concurrent and distributed collection of data by actors in the nodes of a sensor network and analyse the data. Note that the analysis may trigger further data collection, or reuse already collected data. The system being developed also provides a type system for a partial semantic check, to ensure that queries are well-formed, so that for example, a query does not result in a graph of node invocations which is cyclic.

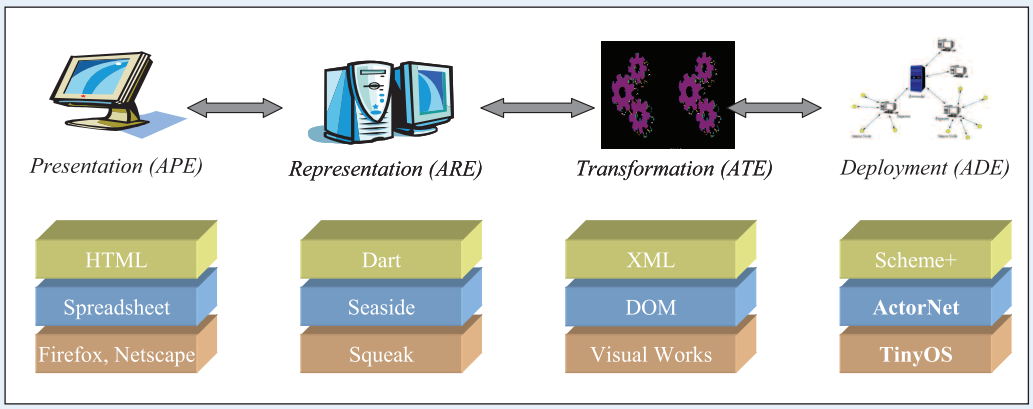

Figure 1: The Ambiance Platform supports macroprogramming WSNs by automated dynamic code generation and deployment.

Ongoing Work

This research is conducted within the framework of the Åmbiance project, funded by the University of Luxembourg, in collaboration with the Open Systems Laboratory (OSL) and the Software Architecture Group (SAG) at the University of Illinois at Urbana Champaign (UIUC-USA), as well as the Computer Science Laboratory (LIP6) of the Pierre and Marie Curie University (Paris).

It draws on our previous realisations Dart, AmItalk, and ActorNet. A first running prototype is implemented and we are finalising the systematic transformation of end-user queries to produce meta-actors and to dynamically manage their life-cycle. Such a lifecycle involves activation and registration, request management, application logic and result dissemination. A type system for Dart is being written in Maude (http://maude.cs.uiuc.edu). Once the core system is in place, we will work with experts in domains to which sensor networks are applicable, such as civil engineering, cooperative target identification and tracking, environment monitoring and security, to experimentally assess and refine the Ambiance architecture.

Links:

http://ambiance.uni.lu/

http://osl.cs.uiuc.edu/

http://wiki.cs.uiuc.edu/SAG

Please contact:

Reza Razavi, University of Luxembourg / FNR, Luxembourg

E-mail: razavi@acm.org

Towards a Platform for Widespread Embedded Intelligence

by Simon Dobson, Steve Neely, Graeme Stevenson, Lorcan Coyle and Paddy Nixon

The vision of pervasive computing is that objects, buildings and environments may be endowed with software intelligence to improve human interactions both with the individual objects and with the system as a whole. Realising this dream is posing significant challenges for designers - allowing individual applications to co-exist in a common space without interfering, making the capabilities of new sensors and services available to other applications as they appear and providing the necessary decentralised control to obtain robust behaviour.

Any large-scale pervasive deployment will face a common core of issues. Decentralised control can be used to provide coordinated behaviour in the face of localised failures and a dynamic population of devices. In many cases a knowledge-based approach has proved to be successful. This is something that is especially important for sensor-driven systems, which are inherently noisy and incomplete and so must deal with substantial amounts of uncertain reasoning. The systems must be self-configuring, -managing and -optimising (the 'self-' properties of autonomic computing and communications) as well as being extensible, standards-based, and upgradable on-the-fly. Finally, systems must scale to hundreds or thousands of devices while offering predictable and robust levels of service.

Within University College Dublin's Systems Research Group we have been looking at middleware solutions to these issues, building both on our own experience with middleware and the other successful pervasive computing frameworks such as the Context Toolkit, Cooltown and Aura. The resulting system - Construct - provides a scalable and open platform on which to build pervasive and adaptive applications.

Construct is unusual in a number of ways. It is entirely data-driven, in that all data is treated as being sensed either by a traditional physical sensor, a 'virtual' sensor accessing on-line information, or a reasoning strategy, making inferences from other data. All information is represented internally using the Resource Description Framework from the semantic web, allowing easy integration with other tools. Applications - both clients' and sensors' - are written to query and/or populate this model, which considerably simplifies the design and deployment of new components.

Noise and uncertainty are ubiquitous characteristics of any sensor-driven system, and Construct applies uncertain reasoning techniques to all its information. The system can maintain consensus estimates of items of interest, modifying its results and confidences over time. This means that the results from excessively noisy or inaccurate sensors, which would tend to disagree with the consensus, will be de-emphasised over time, leading to a self-stabilising infrastructure. The disadvantage of this approach is that all inferences are explicitly tagged with confidence values, which can complicate a programmer's task.

Pervasive systems suffer from frequent node failures, either from genuine faults or from devices being removed or deactivated. This causes problems for many infrastructures: where should data be stored so as to maximise availability? Construct answers this by having each node store a copy of the public state of the system, and gossiping this state randomly between nodes. Any information sensed at a node, for example, may be placed into the public store and will be gossiped to another node. If the source then fails, its observations will have already been transferred to other nodes and so need not be lost. Careful handling is required to ensure both the information is not lost and that the public store does not grow infinitely. Gossiping is extremely robust as it relies on very little knowledge of the network's topology, and makes very fair use of available communications resources without generating hot-spots at popular producers or consumers.

The goal of Construct is to provide a simple framework for the community of pervasive and adaptive systems developers. The platform addresses decentralisation, robustness, fusion and other problems at an infrastructural level, allowing the developer of an individual sensor or service to focus on their own novel contribution. We have developed a number of sensors, including interfaces to the PlaceLab and Ubisense location systems and virtual sensors for a number of web sites, services and standard formats such as iCal. We are actively soliciting the help of others to improve the platform, its services and algorithms.

Construct is being licensed under LGPL. Our hope is that it will serve as a community basis for rapid exploration and integration of new pervasive, adaptive and sensor-driven embedded intelligence.

This work is partially supported by Science Foundation Ireland under grant numbers 04/RPI/1544, "Secure and Predictable Pervasive Computing" and 05/RFP/CMS0062, "Towards a Semantics of Pervasive Computing".

Links:

http://www.construct-infrastructure.org

http://www.ucd.ie/csi

Please contact:

Steve Neely, University College Dublin / IUA, Ireland

E-mail: steve.neely@ucd.ie

Networking Semantic Services for Pervasive Computing

by Valerie Issarny, Nikolaos Georgantas and Sonia Ben Mokhtar

With computing and communication capabilities now embedded in more and more physical objects in the surrounding environment and with most users carrying wireless computing devices, the pervasive computing vision pioneered by Mark Weiser is being realised more than ever. Nevertheless, such an abundance of ambient capabilities in the open, dynamic pervasive environment naturally comes with high heterogeneity, which hinders their ad hoc integration into rich pervasive applications. Within the INRIA ARLES team, we are investigating efficient solutions to pervasive heterogeneity building upon semantic service-oriented architectures.

Various pervasive applications may be provided in today's networked environments. Figure 1 illustrates exploitation of the networked devices to create smart spaces for users. In such an environment, the nomadic user may, in particular, benefit from a large display nearby to watch multimedia content provided by a local media server. The content will automatically be switched to the user's smart phone display as soon as other users get close, to preserve the user's privacy. This application may further be extended to share experiences with others. For instance, assume that the user is actually in a holiday resort, watching an advertisement for available excursions. The user may wish to get feedback about the excursion, which may be obtained through blogs supplied by other nearby visitors on their handhelds and/or the Internet. The user may even wish to look at alternative excursions, whether advertised at nearby resorts, at tourist information offices or even on the Internet. Actuator resources such as displays and software resources such as blogs, multimedia streams and advertisements on the Web, all constitute networked capabilities that may conveniently be combined to dynamically create a pervasive application providing a holiday planner for nomadic users.

Pervasive computing environments, like the one shown above, embed networked devices, possibly wireless, from various application domains, eg, mobile and personal computing, consumer electronics, and home automation domains. This diversity of devices, along with those of networks and of software infrastructures, call for computing paradigms that enable loose, dynamic coupling among heterogeneous resources. Service-oriented architectures (SOA) have emerged as an appropriate computing paradigm resolving a number of heterogeneity issues. However, pervasive environments have further introduced new challenges for SOA both at middleware and at application layer. Regarding the former, devices need to dynamically detect services available in the open networked environment and to interact with them. This requires enabling interoperability of both discovery and communication protocols, as services are implemented on top of diverse middleware platforms (eg, UPnP used in the home and Jini in the mobile domain). Regarding the latter, applications need to compose, in an ad hoc fashion, services available at the specific time and place, which have been developed independently without a priori mutual knowledge. The major challenges are enabling common understanding and adaptation of the semantics of the services being associated and interoperability between service behaviour protocols. A significant enabler for dealing with these issues is the Semantic Web and in particular the Semantic Web Services paradigms, which allow description of service semantics and behaviour and rigorous reasoning about the former. Tackling both semantics and behaviour heterogeneity in a combined way is still an open, challenging issue, further exacerbated by the fact that semantic reasoning is a compute-intensive task, too demanding for the wireless, resource-constrained devices of the pervasive environment.

As part of the work of the INRIA ARLES team, on the development of distributed systems enabling the pervasive computing vision, we are devising enablers for pervasive services. We have, in particular, investigated interoperability enablers and their integration within the base SOA architectural style. The resulting SOA style for pervasive computing offers two key functionalities, each one vertically integrating related advanced features in the middleware and in the application layer, as depicted in Figure 2. First, semantic-rich interoperable service discovery enables the locating of networked services according to their semantic descriptions. In the middleware layer, discovery is achieved, despite the possible coexistence of various discovery protocols, by interworking between them, which further includes middleware data type

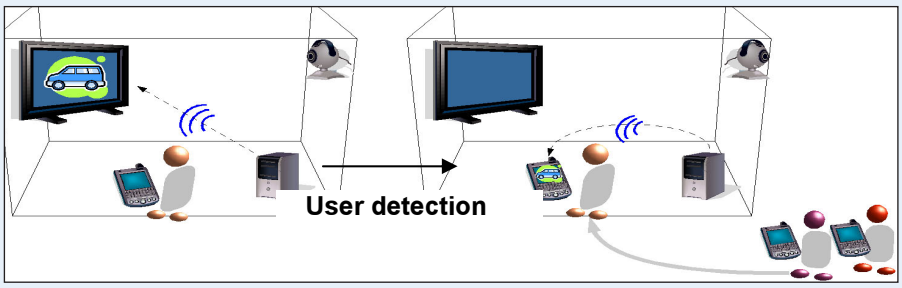

Figure 1: Integrating resources in user-centric smart spaces.

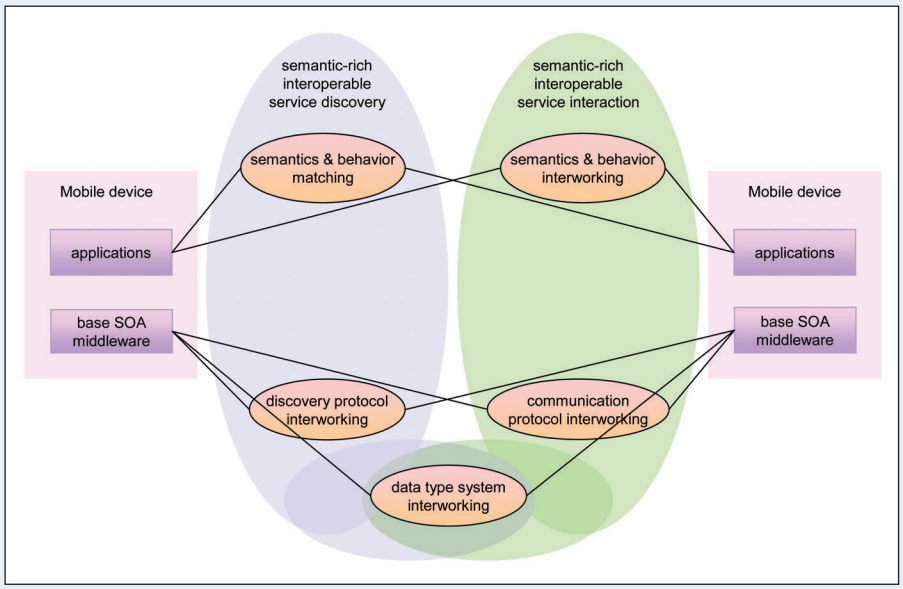

Figure 2: Extending SOA with interoperability enablers.

system interworking. In the application layer, semantics and behaviours of services are reasoned upon and matched with the requested ones in order to decide their adequacy for a given user request. Second, semantic-rich interoperable service interaction enables services to cooperate and compose. In the middleware layer, support is provided for communication protocol and data type system interworking. In the application layer, semantics and behaviours of services are adapted and interworked to enable their seamless interaction. Prototypes of the related architectural components have been implemented. Performance evaluation shows that our solutions comply with the requirements of the pervasive computing environment. We are now in the process of integrating them to offer an instance of the overall architecture. We are further studying extensions to the architecture to effectively exploit beyond third generation networks and related multi-radio wireless handheld devices; this will expand the scope of our interoperability solution to further cover radio and network heterogeneity.

Enabling the pervasive computing vision has given rise to tremendous research over the last decade. In particular, various supporting software infrastructures have been proposed, some enabling semantic awareness for services, others dealing with middleware interoperability. Our contribution lies in the integrated treatment of service interoperability, from the application to the middleware layer. Our solution has further been designed to be lightweight and thus can be supported in most pervasive environments, not requiring the presence of a dedicated server. As a result, true open, pervasive computing environments are enabled - integrating and composing in an ad hoc fashion the most appropriate networked services in order to deliver rich applications to users, independently of heterogeneous underlying software technologies and semantic representations.

Our research is conducted as part of the European IST Amigo and Plastic projects and industrial collaborations with Alcatel and Thales.

Links:
ARLES site: http://www-rocq.inria.fr/arles/

Please contact:
Valerie Issarny, Nikolaos Georgantas,
INRIA, France
Tel: +33 1 39 63 57 89
E-mail: Valerie.Issarny@inria.fr,
Nikolaos.Georgantas@inria.fr

Materials with Intelligence

by Simon Dobson and Kieran Delaney

Sensor networks are the key enabling technology for building systems that adapt autonomously to their environment, without direct human intervention. Most sensor networks operate in free air, but research being conducted in Ireland, between the School of Computer Science and Informatics at UCD Dublin and the Centre for Adaptive Wireless Systems at Cork IT, is starting to explore the tools and techniques we need in order to build 'augmented materials' which combine sensing, actuation and processing into the fabric of built objects.

Embedding sensing into a physical substrate has a number of attractions. Each sensor package can sense a number of local variables such as the stress on the material, its orientation in space, its proximity to other materials etc. Combine these sensors into a network and we can construct a global view of the material and its relationships to the real world. Add processing and we have the potential to build materials that "know themselves", in some sense, and which can react in ways that are far more sophisticated than are possible with simpler, 'smart' materials.

If this all sounds a bit abstract, imagine a person with a broken leg who is wearing a plaster cast. For a physiotherapist, the challenge is to make the person take adequate exercise to stimulate the break, while at the same time stopping them from attempting too much and risking further damage. As the physiotherapy programme changes over time and in conjunction with on-going assessment of

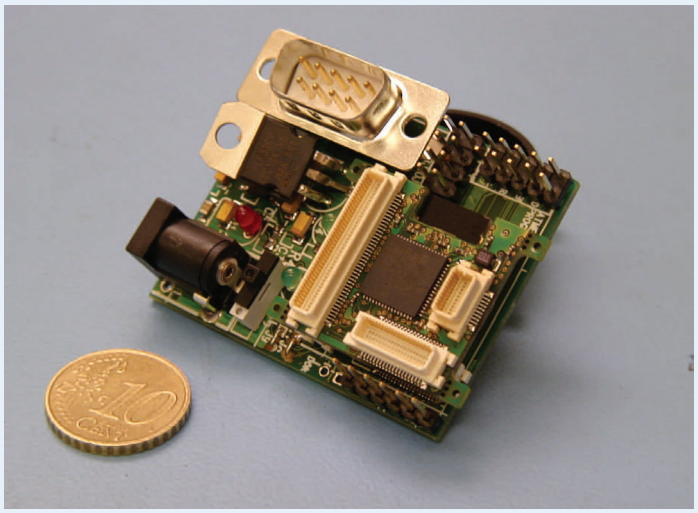

A 2.5cm-on-a-side 'mote', developed by Ireland's Tyndall National Institute.

the injury, the exercise required of the patient and the optimal levels of rigidity and support required of the cast, will also change. If we embed sensing and processing into the cast, the network can sense the load being placed on the cast as the person walks around. This can then be compared with a downloaded therapy programme and react, for example, by glowing green when things are fine, but flashing red lights if the person is overdoing their exercise. It is even possible to build materials with variable rigidity so that the cast adapts the support it provides over the course of treatment.

The individual elements of an augmented material can be based around any 'mote' technology. We are using the platform developed by Ireland's Tyndall National Institute, made available via that institute's National Access Programme. While current mote systems are too large for use in practice, Tyndall's 2.5cm-on-a-side motes (see figure) are being reduced to a 1cm form and beyond, making them realistic for embedded use.

Since a single object might contain hundreds of elements, the elements themselves need to be substantially configured, by making connections to neighbouring elements during curing, for example. Changes to this network may come from node failures, but may also come from physically significant events such as cutting which manifests itself as a (rather structured) partitioning of the network. Dealing with these changes in a tractable way means developing a programming model that operates at the level of the complete material rather than at the level of the individual elements within it. It should also be able to handle failures and errors in sensing in the individual elements. This is not something that is easy to do in traditional languages and we are investigating some techniques pioneered in high-performance computing (skeletons and categorical data types) as a possible basis for building self-managing applications on top of the underlying unreliable sensors and communications.

Augmented materials are in many ways the ideal co-design challenge. The properties of the material determine directly what can be sensed and processed, while software provides capabilities to enhance and complement the material's underlying physics. A physical phenomenon, such as placing one augmented object on top of another, gives rise to individual sensor readings affecting pressure, orientation and the establishing of new wireless communications links etc. These in turn give rise to a semantic inference that can be used in software to drive high-level responses based on the intention inferred from performing this particular action with these particular objects.

While this work is still at a very early stage, we hope that the programme will lead to useful advances in embedded sensor networks, to new forms of microsensing and actuation associated with explicit software control, to improved autonomic control of communications and routing, and to generally useful programming models for sensor-rich multi-agent networks and environments.

Links:

http://www.ucd.ie/csi

http://www.aws.cit.ie

http://www.tyndall.ie/research/mai-group/25cube_mai.html

Please contact:

Simon Dobson, University College Dublin / IUA, Ireland

E-mail: simon.dobson@ucd.ie

The CRUISE Project - Network Initiative for Creating Ubiquitos Intelligent Sensing Environments

by Gabriele Kotsis, Anelia Mitseva and Neeli R. Prasad

Recent developments in research and technology have drawn attention to wireless sensor networks as one of the key enabling technologies of the next ten years. ubiquitous intelligent sensing environments have a promising future in supporting the everyday life of European citizens, bringing important social benefits for each person and for society as a whole. Taking into account the current fragmentation of European research in this field, the CRUISE Network of Excellence (NoE) aims to be a focal point in the coordination of research on communication and application aspects of wireless sensor networking in Europe.

CRUISE's NoE partners are combining their efforts to make a significant contribution to both the coordination and the effectiveness of research into the communication and application aspects of WSN in Europe, promoting not only open cooperation between individuals and organisations from European academia, but also research into cooperation with industry. The main aim for CRUISE is to push for a broad integration of actions and to become one of the focal points for these actions, with the strategic goal of greatly strengthening the European position in this area of research. CRUISE is to amalgamate the majority of European research entities in the field of Wireless and Sensor Networks. It is doing this on the understanding that this critical mass will put the European Wireless and Sensor Networks research community in a position to successfully compete with the US and Asia and to provide support for European industries in this field. These high-level objectives are pursued in activities such as:

- evaluating state-of-the-art wireless sensor networking and collating research topics covered in CRUISE and in Europe into a knowledge model
- depicting a path from current technological status to a successful near-term and long-term vision, by creating the vision and roadmap on necessary research in wireless sensor networking in Europe
- implementing a set of well-defined joint research activities and publishing white papers and a joint book providing an integrated research view
- sharing and integrating methods, software tools and test-beds and making an inventory of the existing platforms and evaluating of the experiences using them
- facilitating the mobility of researchers and PhD students as a key factor for integration and rationalisation of research efforts
- creating distributed PhD teaching curricula and training curricula in sensor networking and related topics; Organising regular CRUISE Schools as a tool for training young researchers and communicating with the industry standardisation community and international research community
- organising open collaboration events with industry, SMEs, academia, other projects with participation of experts from Europe, the USA and Asia, ie, cluster meetings, workshops, roundtable discussions, web-based tutorials, discussion forums, etc., to exchange and discuss ideas and to further elaborate thoughts exploiting CRUISE results.

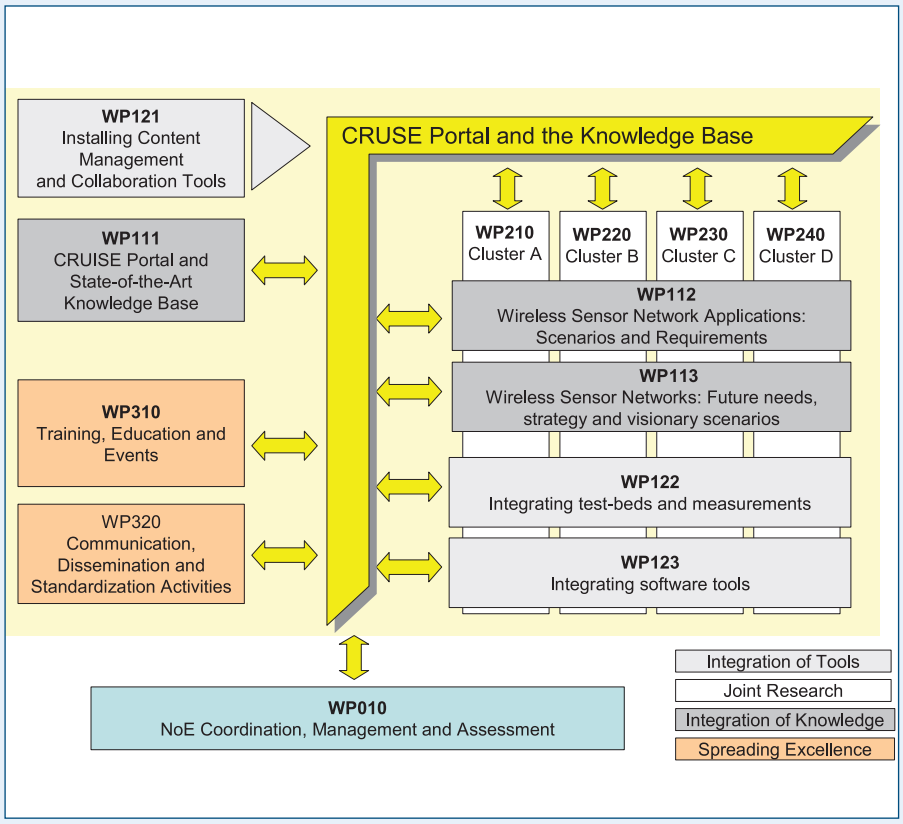

Network integration and research activities.

The joint programme of research activities (see figure), is structured in four self-contained but still inter-related clusters: Cluster A - Architecture and Topology; Cluster B - Protocols and data fusion;

Cluster C - Security and Mobility; Cluster D – Transmission.

CRUISE Services to Industries, SMEs and the Public

One of the main goals of CRUISE is to establish efficient and durable links with industry, using existing relationships and collaboration frameworks as a lever. Each CRUISE partner is to further develop existing liaisons with industry. In order to provide mechanisms of knowledge transfer, the project consortium is to offer Business Deck@CRUISE via the project website. Members of 'Business Deck' are offered access to some of the internal documents prepared by the network. They could influence the choice of topics for tutorials prepared by the network and, to some extent, over topics presented at training schools organised by the partners. They could be also involved in the tools' integration activities.

Cooperation with industry is also to be carried out through links with other IST R&D projects working on close topics, such as eSENSE IP, MAGNET BEYOND IP, GoodFood IP, and through national projects. Special attention will be paid to promoting collaboration with SMEs, which are expected to play a particularly important role in the deployment of sensor networks, but which often lack specific research recourses in order to complement their own skills. Via the CRUISE Portal, SMEs will be able to reach a working group or an appropriate expert from the project with a specific problem to be solved. Adding to the above, CRUISE is to create links with other related IST projects, such as NoE, e-NEXT, NoE NEWCOM, NoE ACE, Coordinated action Embedded WiSeNets.

Last but not least, as part of the plan for using and disseminating the knowledge and to raise public awareness about the immense social benefits of the sensor network applications, active collaboration will be pursued with professional communities and organisations whose interests span different aspects and scenarios of sensing environments apart from the purely technical ones. This will serve the European community beyond academia and science.

This work has been carried out within the framework of the IST-4-027738 NoE CRUISE, which is partly funded by the European Union. The project started on January 1 2006 and has a duration of 24 months. Thirty two internationally recognised groups from Europe, including the ERCIM members VTT and Fraunhofer-FOKUS, are coordinated by the Center for Telenfrastruktur (CTIF) at Aalborg University, Denmark.

Link:
http://www.ist-cruise.eu

Please contact:
Neeli R. Prasad
Aalborg University, Denmark
E-mail: np@kom.aau.dk

IntelliSense RFID - An RFID Platform for Ambient Intelligence with Sensors Integration Capabilities

by Ovidiu Vermesan, Nadine Pesonen, Cristina Rusu, Aarne Oja, Peter Enoksson and Helge Rustad

IntelliSense RFID project expands the technical and market potential of RFID technology by developing multi-protocol devices with sensing capabilities that are able to sense the environment and communicate at different frequency bands.

Embedded intelligence, or ambient intelligence, is the vision of the future where information will be available to people without their being aware of the surrounding information technology itself. In the ambient intelligent world and in the ubiquitous information society, computing is distributed all around in the environment. Users are mobile and communication is set up between individuals, between individuals and objects and between objects. This communication is obtained via 'invisible intelligent devices' that will tag almost everything in our society in an easy and friendly way. These devices will receive, sense, carry, and distribute context information in our environment.

Radio Frequency IDentification (RFID) is the key technology for local connectivity and an early vehicle for readying these future proactive computing systems.

Nanoelectronics allows the mass production of such new silicon RFID devices with sensing capabilities and this makes massive and low-cost tagging of objects feasible.

In this context, the aim of the IntelliSense RFID project is to develop multi-protocol RFID devices with built-in sensing capabilities, operating at multi-frequency bands. The device can be associated with an object, a person or a location through a simple ID and is capable of measuring and acquiring data about the user's behaviour and his environment (such as temperature T, pressure P, humidity H, and pH), thus creating a smart environment based on surrounding 'invisible intelligent devices'.

It is a two-year project (2006-2007) and is part of the Nordic research programme NORDITE. It involves SINTEF (Norway) as the institute

Multi-protocol RFID devices with built-in sensing capabilities used for logistics and mobile phone-based consumer applications.

leading the project, VTT (Finland), Chalmers University of Technology (Sweden), and IMEGO AB (Sweden). The NORDITE programme aims at supporting research institutes and universities from Sweden, Norway and Finland to enhance state-of-the art research in the fields of SW radio, wireless sensors, short-range wireless networks and RFID or MEMS utilising RF technology. The programme goal is to stimulate collaboration among national and Nordic countries, create synergism, and consolidate technological and economic development opportunities for companies and industries. These technology development research projects are funded by TEKES (Finland), VINNOVA (Sweden) and the Norwegian Research Council (Norway).

One functional area of great relevance to many supply chain applications is the ability to monitor environmental or status parameters using a RFID device with built-in sensor capabilities. A generic sensor interface is needed in the RFID platform to enable various applications that can be realised by integrating the proper, low-cost sensor element to the RFID platform. The IntelliSense RFID project focuses on developing a generic sensor interface, integrated with the RFID circuitry, for external capacitive sensor elements. At this stage, capacitive sensors will be used to measure temperature, pressure, humidity, and pH.

One of the major issues with large-scale RFID networks is the lack of standardisation across many fronts. This ranges from the different data formats used, via reduced interoperability between RFID readers and tag/transponder devices from different vendors, to interference problems between RFID products from different manufacturers and the different frequency bands used in different countries.

To overcome such problems, the 'global sensing RFID devices' developed in the IntelliSense RFID project are designed to be able to communicate at different frequency bands with different communication protocols. They can thus communicate with the existing infrastructure (ie, existing RFID interrogators) in various environments. Of special importance is the combination of high frequency (HF) 13.56MHz and ultra high frequency (UHF) 860-960MHz RFID standards. Indeed, multi-band tags are needed to penetrate from logistics to mobile phone-based consumer markets. For example, the same dual-band tag can be exploited for logistics, by using the UHF RFID interrogators, and by consumers, to retrieve the content of the device memory by utilising the integrated HF RFID reader in a mobile phone. Applications that can benefit from such a dual-band RFID tag technology include object authentication and information, personalisation, protection against theft, maintenance, and recycling.

The IntelliSense RFID technology platform will address the HF, UHF and microwave (MW) 2.45GHz frequency bands focusing, in the first phase, on the integration of the HF and UHF RFID specific protocols.

The whole infrastructure 'readers-RFID devices' will form the future 'Internet of Things' that will be one of the main element of the 'Ambient Intelligence' and 'Ubiquitous Information' society.

Please contact:
Ovidiu Vermesan, SINTEF ICT, Norway
Tel: +47 22 06 75 06
E-mail: Ovidiu.Vermesan@sintef.no

Collaborative Capture: A New Perspective for Sensor Networks

by Paul Couderc

Ubiquitous and collaborative capture unveils new perspectives in ambient computing. In the same way that distributed computing proposed approaches to the distribution of data and processing on a set of computing nodes, collaborative capture will have to propose efficient and practical solutions to recompose structured data from items collected asynchronously from contributing nodes.

An important aspect of ambient intelligence is to link information systems with real work. To this end, sensors and networks of sensors play an important role in providing data and events from the real world. Sensors and networks of sensors are thus gaining more and more attention from the research community. Enabling cooperation between these tiny devices, in an extremely dynamic context, poses real challenges. In particular, we have to cope with very limited resources (energy, processing power, memory), high failure rate and node mobility. Another important issue is to offer programming abstractions and system architectures to virtualise a collection of sensors such that it can be used as a single virtual resource. The main application domain for these devices is monitoring at large (environmental, medical, industrial, etc.)

Surprisingly, much less attention has been paid to what is becoming the most important sensor network ever deployed: the multimedia phone. While these are well-known as personal messaging devices, their role as sensors has been largely ignored, in spite of their powerful capabilities in this respect. These devices can capture sound, picture, video, can cooperate with other sensors such as GPS positionning receivers and, of course, can communicate with global as well as local networks.

An important challenge for these existing and promising sensing infrastructures is to explore their potential applications and to address the problems raised by their emergence as an unexpected sensor network.

The ACES (ambient computing and embedded systems) research team at INRIA has a strong background in the domain of cooperating objects, which means groups of devices which spontaneously team up to form a short-lived information system (called spontaneous information system) in order to accomplish a task. It has demonstrated several applications of this concept. Recently, the team proposed the concept of spontaneous hotspots. This is where swarms of mobile phones are used to increase download-efficiency of context-related resources in a GPRS network using bluetooth-enabled cooperation (Spontaneous Hotspots: Sharing Context-Dependant Resources Using Bluetooth, Xavier Le Bourdon, Paul Couderc, Michel Banâtre International Conference on Networking and Services (ICNS'2006) July 16-19, 2006 - Santa Clara, CA, USA). This approach both balances the network download cost between users and reduces the network load, as well as decreasing download times. A typical usage scenario would be, for example, at a sports event, where many people in a stadium would like to see an instant replay of the recent action. In this context, the same resources (video of the replay) are requested by many users. The spontaneous hotspot allows the network to balance the download and to flood the replay into the stadium through bluetooth.

We are now investigating a related, but more innovative concept: collaborative capture, where the shared data is no longer resources hosted in the network, as in the previous case, but structured information composed of elements contributed dynamically by a collection of nodes. Consider, for example, someone arriving late in a meeting. Many phones, treated as audio capture devices, could have recorded the beginning of the meeting to allow the newcomer to catch up with the missing part. Similar scenarios with photos can be envisaged, allowing for example, a group of people

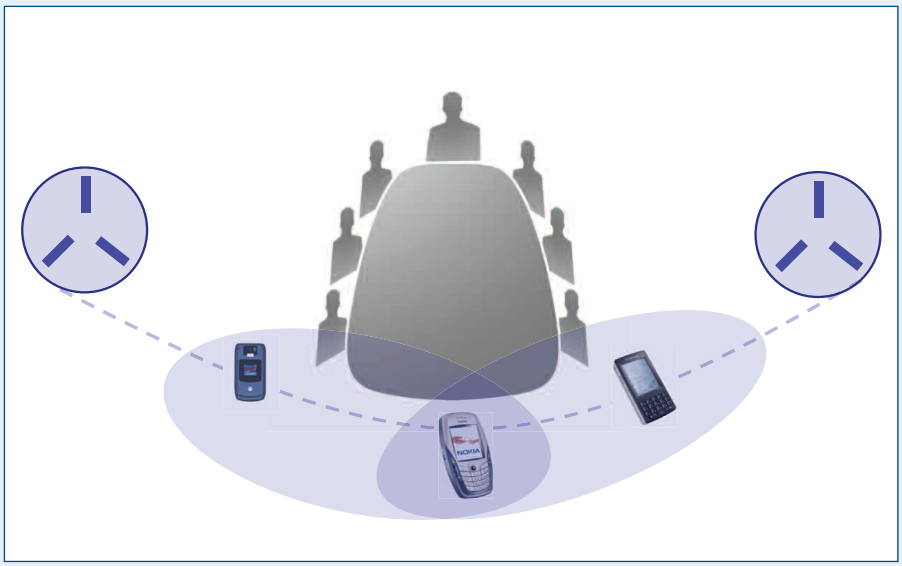

Many phones, treated as audio capture devices, could have recorded the beginning of the meeting to allow the newcomer to catch up with the missing part.

taking pictures at an event to merge their captures and provide a complete collection. Another interesting use of collaborative capture is to use several nodes to increase the quality or the accuracy of a given data, by combining their individual data.

These applications, which can appear simple, raised difficult issues. Context descriptions of available resources and on-going capture processes should be accurate enough to allow sharing, without overloading the local communication medium to allow their efficient distribution on-the-field. Obviously, privacy is a big issue. For example, you may agree to share the pictures you are taking at a given event with only a few of the people present, but not the pictures taken an hour ago in another context. Traditional security, with authorisations, is difficult to apply in practice in the context of spontaneous interactions with new and potentially unknown nodes. Innovative approaches, based for example on more dynamic notions such as trust and behaviour, have to be investigated before practical security mechanism for these systems can be found. In relation to trust, we must be able to rank the quality of the contributions, which may include poor data, or even false data. Worse, the very notion of ubiquitous capture can be frightening: the potential capture activity of anyone, anywhere may change social relations between people. It's too late to do anything about that, this distributed infrastructure is already in our pockets!

Link:
http://www.irisa.fr/aces/

Please contact:
Paul Couderc, INRIA-IRISA, France
E-mail: paul.couderc@inria.fr

Autarkic Power Generation for Networked Systems

by Eberhard Mueller and Ulrich Hofman

A research project funded by the FIT-IT programme of the Austrian ministry for traffic, innovation and technology (BMVIT) investigates the potential of different energy sources for communication purposes, especially for out-door activities.

The economic relevance of this research activity can easily be seen by examining the efforts companies take to optimise the energy consumption of embedded devices. The use of alternative energy supply, like Piezo-Elements or Thermo-Elements, will thus be a great chance to achieve this goal so as to minimise costs and to make devices independent from external energy supply.

Energy management is not only considered in terms of minimising the energy consumption, but also as finding additional ways for energy to be supplied. Furthermore, the energy-consuming activities are bound to the current status of energy production and the current cost functions. For example, energy consuming communication tasks are carried out not continuously, but at optimal opportunities by episodic transmission protocols. This principle is verified using two very different application scenarios: skis and bicycles.

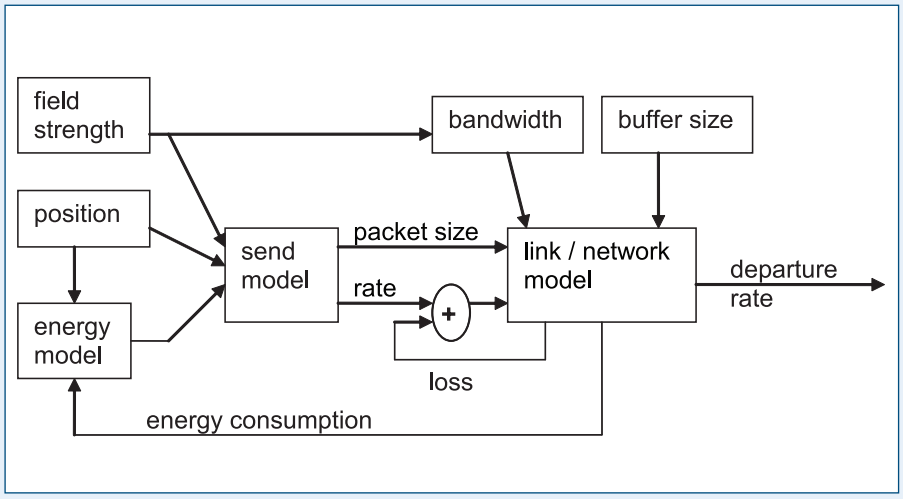

Figure 1: Logical components of Unified Model.

The scientific aim of the project is to develop an innovative integrated framework and tool for modelling the energy and communication processes in networked embedded systems. Today's modelling approaches and tools are based on a combination of discreet modelling of the communication processes and an underlying model describing the energy consumption. The modelling and simulation of the energy processes in sinks (sensors, actuators, physical communication) and sources (generator, battery) are realised by means of continuous models.

This two-paradigm approach has its disadvantages:

- Handling two modelling approaches and isolated tools is difficult for the researchers/engineers and increases the duration of the development cycle.
- Interfaces between different modelling approaches/tools have a potential for misinterpretation by the experts in their respective domain.

To overcome these drawbacks we model the data communication process by following a new approach that we have investigated in the European IST project INTERMON. There, data traffic is modelled as a continuous (in time and value), fluid process. We bring this paradigm into the modelling domain of networked embedded systems. Using the fluid abstraction, which was successfully employed in large scale IP networks, will open the way to a unified continuous modelling framework. This framework will accommodate the energy and communication models that will be developed within the project. In a second stage, the models will be validated and integrated into a complex system. As a proof of concept this new integrated modelling approach will be employed for the design and the development of a prototype.

The main challenge is to optimise data communication, sensor data generation and storage of relevant data in the prototype. This includes an autonomous power supply which makes use of the forces, temperatures or other physical phenomena that occur during the use of such devices. State-of-the-art sensor technologies, power sources and communication systems will be investigated and their process models will be derived, measured and integrated into the continuous unified modelling framework. Measurement data of field trials will be used for the validation of the scientific and technical approaches.

Figure 1 shows the logical components and models which are developed. The send model includes continuous and episodic transmission protocols above the Layer-2 WLAN-protocol (Figure 2). Figure 3 shows transfer intervals of a mobile station to different base stations.

Figure 2: Optimal data transfer to a base station.

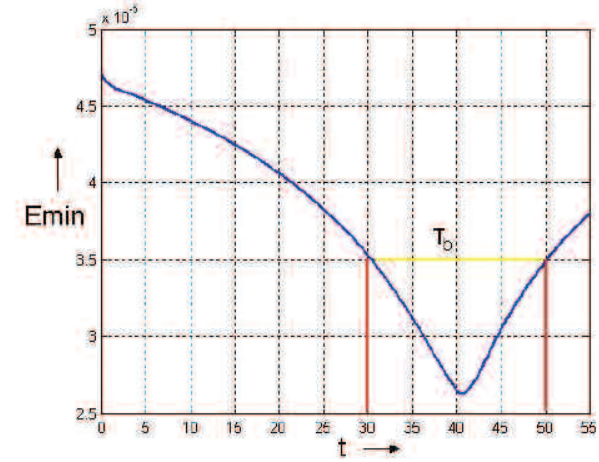

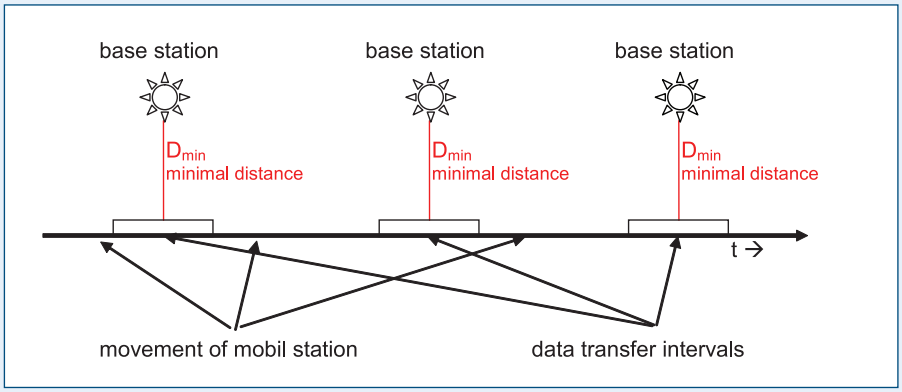

Figure 3: Data transfer to different base stations.

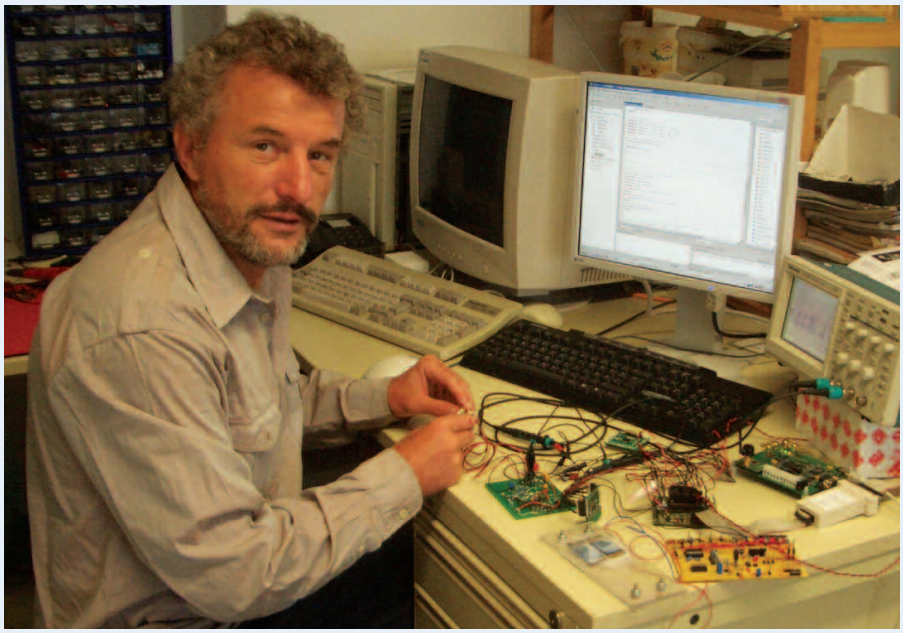

Figure 4: The author working with Piezo-Elements and Embedded System.

Links:

Salzburg Research:
http://www.salzburgresearch.at/company/
ESYCS - Embedded Systems Cluster
Salzburg: http://esycs.salzburgresearch.at/
INTERMON: Advanced architecture for
INTER-domain quality of service
MONitoring, modelling and visualisation:
http://www.ist-intermon.org/

Please contact:
Eberhard Mueller and Ulrich Hofmann
Salzburg Research GmbH, Austria
Tel: +43 662 2288 445
E-mail: emueller@salzburgresearch.at,
ulrich.hofmann@salzburgresearch.at

Formally Bounded Embedded Systems

by Kevin Hammond

The EU-funded EmBounded project aims to develop new techniques for providing formal guarantees on resource usage bounds. It is a collaborative project between AbsInt GmbH, a high-tech company specialising in high integrity real-time systems, and academic partners in the UK, France and Germany.

A major difficulty with constructing embedded systems is that the available system resources (processor, computer memory, power) are necessarily restricted due to cost and other considerations. If it were possible to determine strong bounds on the use of such resources, then potentially there would be significant benefits in terms of manufacturing cost, reliability and performance. However, determining such bounds automatically is a difficult problem and doing so manually is becoming impracticable as embedded software increases in complexity.

The EmBounded project, led by Dr Kevin Hammond at the University of St Andrews, aims to research this problem. Our vision is one where certificates of the bounds on resource usage can be obtained from a source program, through automatic program analysis, independently of the usual software compilation process. These certificates may then be verified using formal proof techniques based on a program logic that captures the meaning and behaviour of programs.

The EmBounded project builds on European strengths in programming language design, program analysis and embedded applications. Our work is based around a new domain-specific programming language notation, Hume, that aims to find an optimum point between good programming abstraction and the ability to derive good cost information. Hume embeds a rule-based notation for sequential computations within a finite-state-machine-based notation for concurrency control. The process notation is designed to be easily tractable to standard analyses such as model-checking for deadlock detection; while the rule-based notation is designed so that we can easily expose information about cost. Hume is designed to be practical as well as having interesting language properties. Prototype implementations of Hume produced at Heriot-Watt University can run in less than 30KB on a standalone Renesas M32C development board, with minimal dynamic memory requirements, depending on the end-application.

***embound**, v.*
poet. arch.
trans. To set bounds to; to confine, contain, hem in.
Hence **embounded** *ppl. a.*
1595 SHAKESPEARE: *The Life and Death of King John* IV. iii 137
That sweete breath which was embounded in this beauteous clay.

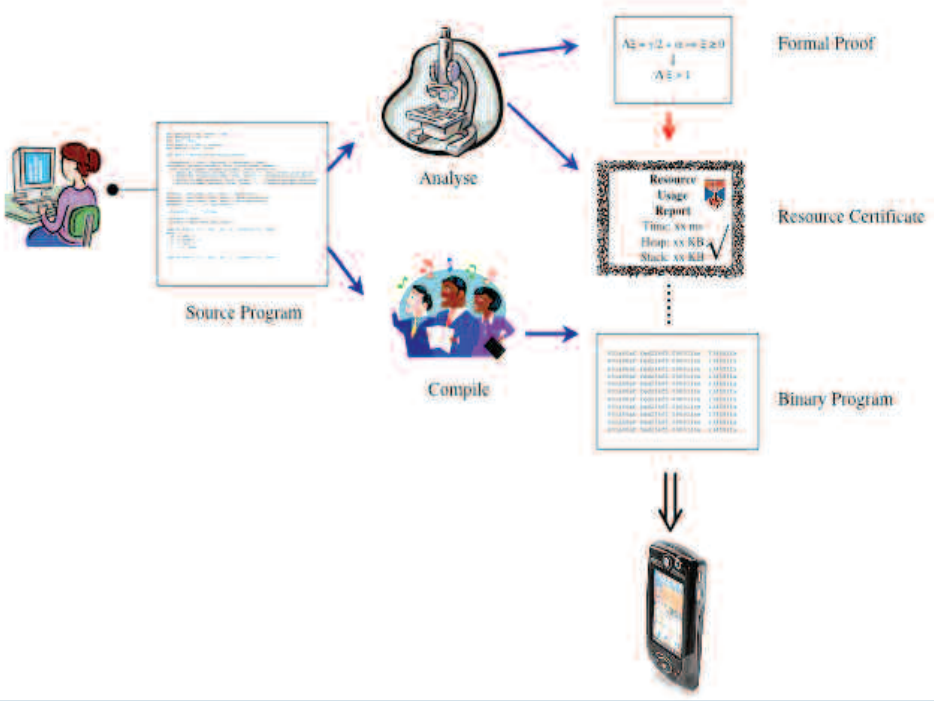

The 'EmBounded' vision.

To date, project researchers at Ludwig-Maximilians-Universität, München, and St Andrews University have developed new theoretical models of bounded resource usage for dynamic memory and time metrics based on Hume program source. They have related these formally to the underlying machine operations using an abstract-machine approach. We are now constructing automatic analyses based on these models that will expose resource constraints and other program properties in our program logic. This will give a good model of program cost based on the program structure.

In order to obtain accurate information about time usage on a specific computer architecture, we are using the aiT tool developed by AbsInt GmbH to provide precise and guaranteed timing costs for fragments of machine code programs. This low-level

information can then be combined with the high-level information obtained from our program source to provide a highly accurate analysis that will give guaranteed bounds on program execution time.

In order to produce convincing results, it is necessary to test our approach using sophisticated and realistic applications. Researchers at the Laboratoire des Sciences et Matériaux d'Électronique — LASMEA (Clermont-Ferrand, France), have been developing a number of real-time computer vision algorithms in Hume. These algorithms can be exploited to detect, for example, road features from a moving vehicle, an essential component of any self-controlled (autonomous) road vehicle. In due course, we intend to apply these algorithms to the sensor and control systems used in the CyCab autonomous vehicle (http://www.robosoft.fr), a self-controlled electric car resembling a golf buggy, and capable of speeds up to a rather alarming 30 km/h.

Links:
Project: http://www.embounded.org
Hume: http://www.hume-lang.org.

Please contact:
Kevin Hammond, University of St Andrews, Scotland, UK
E-mail: kh@dcs.st-and.ac.uk

Adaptable and Context-Aware Trustworthiness Evaluation

by Gabriele Lenzini, Santtu Toivonen and Ilkka Uusitalo

Conducted as part of the EU-ITEA Trust4All project, which aims to enhance dependability in networked and component-based systems, VTT Technical Research Centre of Finland and Dutch Telematica Instituut began investigating trustworthiness evaluation in ubiquitous environments. So far, the research has concerned context-aware trust and trust evaluation with indirect information.

Trust is an increasingly important phenomenon to grasp and support in open environments, such as where networked devices operate. A common scenario sees a Trustor, the subject of trust, search for a Trustee, the object of trust. The Trustor's trustworthiness evaluation directed to the Trustee should be facilitated. Semi-automatic trustworthiness evaluation is of special relevance in embedded systems as well as in the semantic web, or in the ubiquitous environments, where the Trustors and Trustees can be computer programmes in addition to human beings. To perform an appropriate evaluation, Trustors request various information characterising Trustees. Typical solutions ask for users' trust credentials, often expressed in terms of user profile, reputation, and recommendations. The difference between reputation and recommendation here is that reputation is based on the experiences of the Trustor, whereas recommendations are communicated experiences of others.

Contextual data, that is information considered relevant to the interaction between a Trustor and a Trustee, including the environment in which this interaction occurs, can influence the result of the trust establishment. For example, acting in a safe environment can affect the trust that an application has in a component. Context-awareness has been recognised in many research areas of information technology, such as information filtering and retrieval, service provisioning and communication. However, the relationship between context and trust has not received very much attention, apart from some occasional work. This is unfortunate, since such relationship can easily be recognised and its existence justified.

In the first phase of our research, we contextualise trust-evaluation by considering individual context attributes and assign them with values influencing the trustworthiness evaluation process. Depending on the importance of a given context attribute, determined by what we call a trust purpose, weights can be assigned to amplify or weaken the respective attribute's influence on the overall trust. For example, the Trustee's location can have significant importance when deciding if to grant access to the local wireless network. A policy can be introduced so that nobody (regardless of possessing the network key or not) is allowed to use it if residing outside the company boundaries. We apply the notion of context also to reputations and recommendations. When dealing with reputations and recommendations, contextual data can be used to emphasise other experiences that have taken place under "similar enough" conditions to those in which the Trustor currently finds himself. Here "similarity" between contexts can be computed, for example based on reasoning or on context ontologies, and expressed as minimal contextual distance.

The above-mentioned trust evaluation works on the assumption that reputation data and recommendations are directly available to the Trustor. In the second phase of our research, we extended the adaptability of context-aware trustworthiness evaluation by considering scenarios where this direct information is not necessarily available to the Trustor. In such cases, we highlight three subcases to be studied. First, the Trustee's behaviour in other contexts may be unknown. Secondly, the Trustee's behaviour in the current context may be unknown, for example due to absence of knowledge referring to the current context. Thirdly, reliable recommenders and/or recommendations about the Trustee might not be available or they could refer to the Trustee acting in different context. We address each of these

cases with separate proposals, all of which are based on similarities between the entities - Trustor, Trustee, recommender and other entities connected to each other by a (social) network - or contexts attached to the entities. For example, the lack of direct information about a new scientific conference does not stop the Trustor inferring trustworthiness by considering the reputation of the conference chair and favourable recommendations about the programme committee members and the conference proceedings publisher.

Trust4All is a two-year EU-ITEA project. Trust4All aims to enhance dependability and trustworthiness of networked and component-based embedded systems. Trust4All project partners come from industry, research institutes, and academia: CWI, Océ –Technologies, Philips Research (the project coordinator), Telematica Instituut, Univ. Eindhoven (TU/e), Univ. Leiden (the Netherlands); Nokia, Solid Information Technologies, VTT (Finland); ESI, FAGOR, IKERLAN-Electrónica, Robotiker, Visual Tools (Spain).

Links:
Telematica Instituut: http://www.telin.nl/
VTT: http://www.vtt.fi
EU-ITEA project Trust4All:
http://www.win.tue.nl/trust4all/

Please contact:
Gabriele Lenzini, Telematica Instituut,
The Netherlands
E-mail: gabriele.lenzini@telin.nl

Santtu Toivonen, VTT Technical Research Centre of Finland
E-mail: sainttu.toivonen@vtt.fi

Digital Identity in Ambient Environments

by Ben Schouten and Onkar Ambekar

Embedded systems and ambient technology enable users to interact at any time and anywhere. In the BASIS project for identity management, CWI investigates transparent biometrics in home environments. Possible application areas are user profiling for shopping , listening to one's favourite music and operating gadgets and appliances in the home.

Historically, personal relationships, face-to-face encounters, notaries, and third party counsel verified our identity. Nowadays, reliance on paper has shifted to electronic documents and in the same way, reliance on traditional trust factors has had to shift to electronic authentication for exchanging information, goods, and services. As territories and communities are newly defined in the digital age, there is a need for a new definition of digital identity. The question is: Do we possess only one 'federal identity' as in our passport, or do we allow the use of partial identities, depending on the application we are using or on the relationship or community of which we are a part? (See Figure 1.)

The research project Biometric Authentication Supporting Invisible Security (BASIS) tries to answer these questions. It aims to develop authentication protocols in networked and distributed systems and, in particular, studies robust face recognition in images from multiple cameras. Users will have an aware, adaptive and responsive personal space, in the home environment. The aim is to investigate the possibilities of biometric authentication for securing access to information and services in this personal environment, with a focus on user convenience and privacy protection.

The research project addresses three topics in the following work packages (WPs). The WP Transparent Biometrics is about transparent biometric authentication (read without taking any specific action) as a means to enhance user convenience - for instance, to automatically hear your favourite music. The WP Template Protection is about anonymous biometric authentication as a means of protecting the user's privacy and the WP Home Biometrics is about the specific problems of biometric authentication in the home environment. Research at CWI, WP leader for Home Biometrics, started in January 2006.

In the home environment, it is unlikely that biometrics based on the recognition of one modality - such as speech or vision - will suffice. If we want to provide the user with easy and transparent use we should be able to identify subjects in changing and uncontrolled conditions. In the project, identity is not defined as a unique set of personal data, like name or age, but more generally. Height, weight, or anything that can be used to classify user profiles is used. For example, we are not interested in the

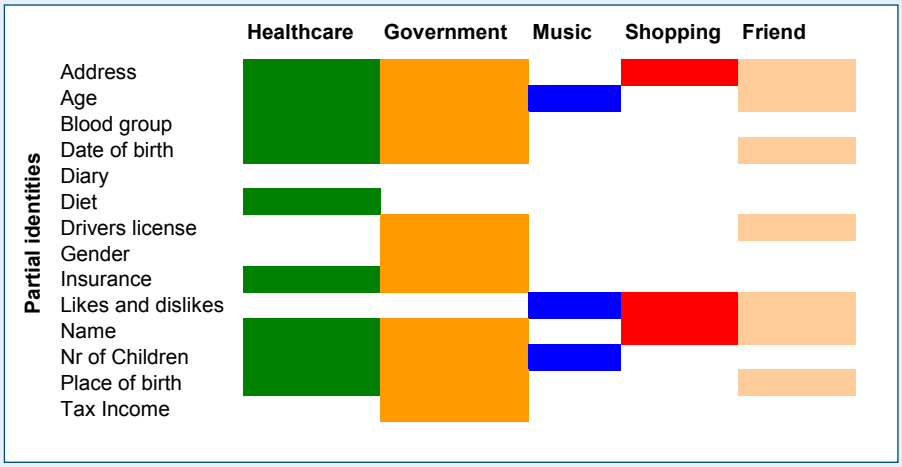

Figure 1: People share different identity data with different organizations and people.

name of the toddler who is approaching a dangerous stairway in order to automatically close the stair gate.

With BASIS, information from different sensors (speech, video etc) will be fused and analysed to produce a coherent interpretation. The main complexity here stems from the heterogeneity of the data sources. As a consequence of small user groups in the home environment it makes sense to include so-called lightweight or 'soft' biometrics like weight, height or the location of a subject to improve reliability.

Up to now, many ways of fusing data have been investigated. In the case of a recognition system, they can roughly be divided into four groups according to the level on which fusion takes place: pre-processing level, feature level, likelihood score level and decision level. As we will cope with uncertain inputs and changing modularity, the best candidate is the likelihood score level. A very promising candidate for solving the fusion task is to use a Bayesian network (BN) or an extension. Using BN, one can cope with uncertain inputs and prior knowledge (in the form of beliefs and statistics), which could be incorporated in a BN. Also Kalman-filters and HMM approaches, used widely and successfully for coping with noisy inputs and tracking, could be expressed in terms of Bayesian networks.

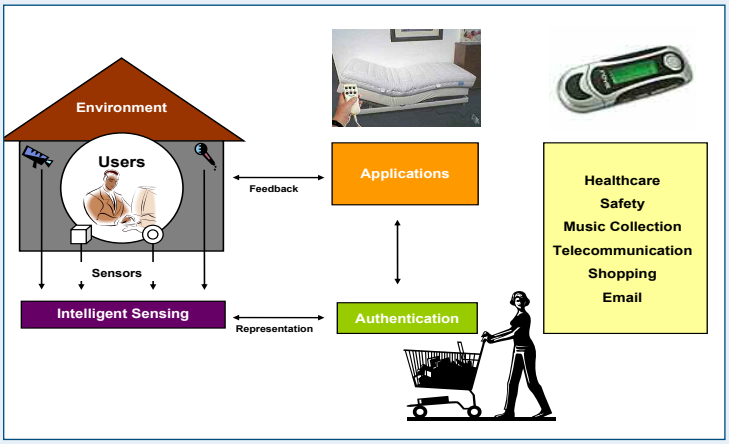

Figure 2: By ambient technology, users can authenticate and profile various applications.

The project outcomes can support future applications in smart environments, adaptive physical environments, profiling applications, and parental restrictions (see Figure 2). In the future we foresee people carrying wireless identity tokens, for example in a handheld phone, RFID or implanted chip, constituting partial identities by which people present themselves, enabling them to authenticate transparently by using different applications. To avoid privacy constraints and repudiation, leaving control in the hand of the user, negotiation and interaction in the authentication process, are essential. Just as in the 'old days' when police would take down your name and address as a consequence of jumping a traffic light.

The research consortium consists of University of Twente, University of Eindhoven, CWI and Philips HomeLab and is co-financed by IOP GenCom of the Dutch Ministery of Economic Affairs. In the course of the four-year project the results will be demonstrated in the Philips Homelab. In the short time the project has been in existence, results have already led to a patent application.

Link:
http://www.cwi.nl/pna4

Please contact:
Ben Schouten
CWI and FONTYS, The Netherlands
Tel: +31 20 5924050
E-mail: B.A.M.Schouten@cwi.nl

Trust4All enhances Trustworthiness of Embedded Component-based Systems

by Gabriele Lenzini, Johan Muskens and Andrew Tokmakoff

As part of the EU-ITEA research programme, Trust4All project partners from industry, research institutes and universities teamed up to design and develop an embedded systems framework that helps component-based applications to maintain a high level of dependability and security. Philips research and the Dutch Telematica Instituut designed and developed the framework's trustworthiness management features.

Component-based Software Engineering is an emerging development paradigm that promises to accelerate software development and to reduce costs by assembling systems from pre-fabricated software components. It has also been adopted for the development of software for embedded devices.

There are several trends in component-based embedded devices. First, systems are becoming more open in the sense that more third party components (ie, software components that have not been developed by the device vendor) are used for their development. So, component-based devices run various components that interact with each other to provide the high-level system functionalities. Furthermore, systems are becoming more dynamic in the sense that the software configuration of a device changes when the device is owned and used by the end user. These trends enable an increasing number of features to be offered by

embedded devices, but threaten the devices' dependability and security properties. In the context of an open component-based system that supports unmanaged upgrading and extension (changing the software configuration), it is possible that component-based devices may host unstable components that could negatively affect the devices' dependability. Here, the most risky scenario is the download of new components to extend the system's capabilities; this exposes the system's internals to potentially malicious components that may compromise the system's overall security.

Recent research indicates that the use of Trust Management can be a useful tool when addressing the dependability and security concerns encountered in distributed and embedded systems.

The Trust4All project has embraced this emerging topic and investigated its application in the context of component-based embedded systems. Our approach has been to extend the component-based middleware, first developed within the EU-ITEA Robocop and Space4U projects, to include a Trustworthiness Management Framework (TMF), as shown in the figure. This framework provides low-level mechanisms that can be used to control and enforce components' behaviour, based upon established Trustor-Trustee relationship. The goal of TMF is to support easy and late (possibly runtime) integration of components and still have dependability and security properties that are satisfactory to the user.

The TMF design addresses the following challenges:

- The TMF acts on behalf of Trustors and explicitly takes their trust requirements into account.
- The TMF makes use of 'metric sources' to monitor and report on the system's overall behaviour.
- The TMF makes use of 'actuators' to control a Trustee's mode of operation, thereby influencing the system's behavioural characteristics.
- The TMF makes use of the Space4U Resource Management Framework to decide whether a component's requested mode of operation should be allowed from a resource consumption point of view.

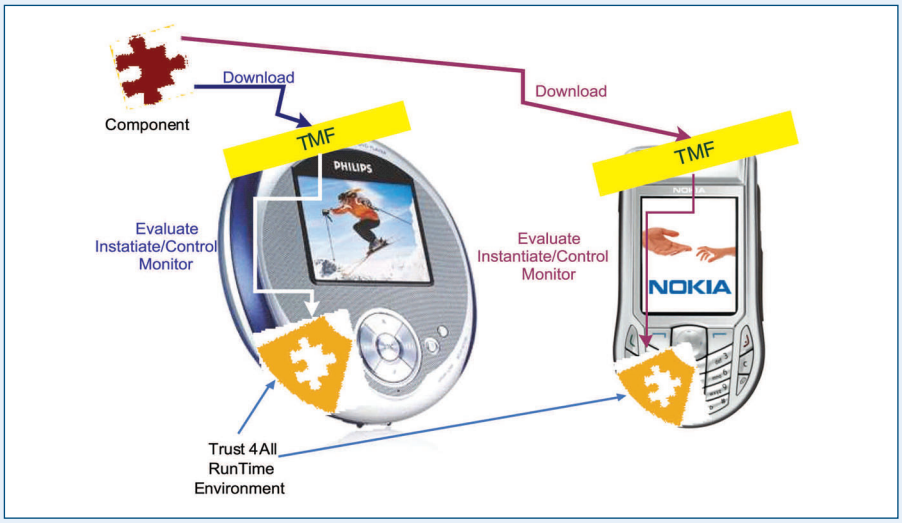

Picture: Telematica Instituut and Philips.

The Trust Management Framework, as part of the middleware of the embedded component-based devices, ensures the dependability of the system, when a new component is added.

A core part of the TMF is the Trustworthiness Evaluation Function (TEF), which is responsible for calculating the trustworthiness of a component/Trustee. It can be parameterised with Trustor-specific dependability and security requirements. By analysing a Trustee's 'estimated quality attributes' and recommendations, the TEF can evaluate a component's 'present trustworthiness' (or compliance) in terms of the degree to which it satisfies the Trustor's stated dependability and security requirements. Through analysis of component's behaviour, the TEF calculates (utilising Subjective Logic) the 'believed trustworthiness' of a component in terms of the following aspects:

- *Benignity:* the belief that the component will continue to satisfy the Trustor's requirements,
- *Stability:* the belief that the Trustee's behavioural qualities of will remain within a given neighbourhood.

The triple of compliance, benignity, and stability is used to make control decisions that regulate the system's overall dependability and security characteristics. For example, when a component with positive compliance begins to show a reducing benignity, the TMF may decide to 'wrap' the component in an attempt to minimise its potential impact on the system, since it is displaying increasingly less-compliant behaviour. Similarly, when there is a strong disbelief in stability, this may indicate a need to re-evaluate the component's estimated quality attributes and, as a consequence, re-evaluate its compliance.

The project began in July 2005 and will run for two years. Its efforts are now moving from research and design towards implementation and construction of demonstrators in the domains of consumer electronics, mobile devices, and domotics.

Many companies and researchers collaborate in Trust4All: CWI, Océ –Technologies, Philips Research (project coordinator), Telematica Instituut, Eindhoven University of Technology (TU/e), Univ. Leiden (the Netherlands). Nokia, Solid Information Technologies, VTT (Finland), ESI, FAGOR, IKERLAN-Electrónica, Robotiker, and Visual Tools (Spain).

Links:

http://www.research.philips.com
http://www.telin.nl
http://www.win.tue.nl/trust4all

Please contact:

Gabriele Lenzini, Telematica Instituut,
The Netherlands
Tel: + 31 53 4850463
E-mail: Gabriele.Lenzini@telin.nl

Johan Muskens, Philips Research Europe,
The Netherlands
Tel: + 31 40 2742491
E-mail: Johan.Muskens@philips.com

Nanotechnology

Single-Molecule Switch and Memory Element

by Emanuel Lörtscher and Heike Riel

Scientists at the IBM Zurich Research Laboratory have demonstrated how a single molecule can be switched between two distinct conductive states, which allows it to store data. As published in SMALL, these experiments show that certain types of molecules reveal intrinsic molecular functionalities that are comparable to devices used in today's semiconductor technology. This finding is a result coming from IBM's Research Labs' effort to explore and develop novel technologies for beyond the CMOS (Complementary Metal Oxide Semi-conductor) era.

In the August 4 issue of SMALL[1], we report on a single-molecule switch and memory element. Using a sophisticated mechanical method, we were able to establish electrical contact with an individual molecule to demonstrate reversible and controllable switching between two distinct conductive states. This investigation is part of our work to explore and characterize molecules to become possible building blocks for future memory and logic applications. With dimensions of a single molecule on the order of one nanometer (one millionth of a millimeter), molecular electronics redefines the ultimate limit of miniaturization far beyond that of today's silicon-based technology.

The results show that these molecules exhibit properties that can be utilized to perform the same logic operations as used in today's information technology. Namely, by applying voltage pulses to the molecule, it can be controllably switched between two distinct 'on' and 'off' states. These correspond to the '0' and '1' states on which data storage is based. Moreover, both conductive states are stable and enable non-destructive read-out of the bit state - a prerequisite for nonvolatile memory operation - which we demonstrated by performing repeated write-read-erase-read cycles. With this single-molecule memory element, we have documented more than 500 switching cycles and switching times in the microsecond range.

Crucial for investigating the inherent properties of molecules is the ability to deal with them individually. To do this, we extended a method called the mechanically controllable break-junction (MCBJ). With this technique, a metallic bridge on an insulating substrate is carefully stretched by mechanical bending. Ultimately the bridge breaks, creating two separate electrodes that possess atomic-sized tips. The gap between the electrodes can be controlled with picometer (one thousandth of a nanometer) accuracy due to the very high transmission ratio of the bending mechanism. In a next step, a solution of the organic molecules is deposited on top of the electrodes. As the junction closes, a molecule capable of chemically bonding to both metallic electrodes can bridge the gap. In this way, an individual molecule is 'caught' between the electrodes, and measurements can be performed.

The molecules investigated are specially designed organic molecules measuring only about 1.5 nanometers in length, approximately one hundredth of a state-of-the-art CMOS element. The molecule was designed and synthesized by Professor James M. Tour and co-workers at Rice University, Houston, USA.

The main advantage of exploiting transport capabilities at the molecular scale is that the fundamental building blocks are much smaller than today's CMOS elements. Furthermore, chemical synthesis produces completely identical molecules, which, in principle, are building blocks with no variability. This allows us to avoid a known problem that CMOS devices face as they are scaled to ever smaller dimensions. In addition, we hope to discover possibly novel, yet unknown properties that silicon and related materials do not have.

[1] (E. Lörtscher, J. W. Ciszek, J. Tour, and H. Riel, "Reversible and Controllable Switching of a Single-Molecule Junction", Small, Volume 2, Issue 8-9 , pp. 973-977, 4 August 2006)

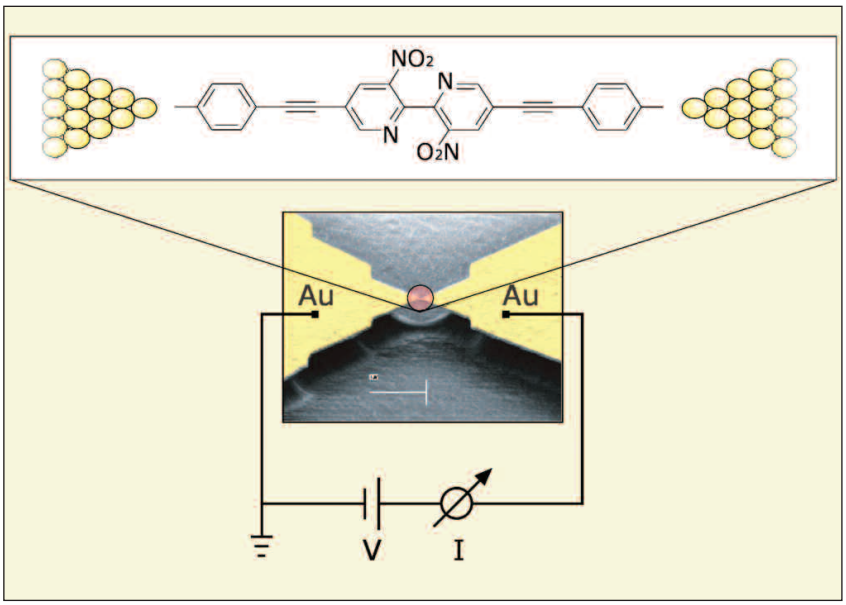

Figure 1:
Contacting a single molecule: Scanning electron microscope image of a metallic bridge, used for establishing contact to an individual molecule. Atomic-sized tips, serving as electrodes, are created by stretching and finally breaking this 'nano' bridge. The method is called the mechanically controllable break junction (MCBJ). The sketch at the top shows the switching molecule "caught" between the electrodes by closing the bridge gradually with picometer accuracy until one molecule "reaches" both electrodes. The switching molecule is a specially designed organic molecule (bipyridyl-dinitro oligophenylene-ethynylene) with dimensions of typically 1.5 nm.

Promising Nanotechnology for Beyond the CMOS Era

The single-molecule switch is part of a series of results achieved by IBM researchers in their efforts to explore and develop novel technologies that will surpass conventional CMOS technology. Miniaturizing the basic building blocks of semiconductor chips, thereby achieving more functionality on the same area, is also referred to as scaling, which is the main principle driving the semiconductor industry. Known as 'Moore's Law', which states that the transistor density of semiconductor chips will double roughly every 18 months, this principle has governed the chip industry for the past 40 years. The result has been the most dramatic and unequaled increase in performance ever known.

However, CMOS technology will reach its ultimate limits in 10 to 15 years. As chip structures, which currently have dimensions of about 40 nm, continue to shrink below the 20 nm mark, ever more complex challenges arise and scaling appears to be no longer economically feasible. Below 10 nm, the fundamental physical limits of CMOS technology will be reached. Therefore, novel concepts are needed.

Link:
http://www3.interscience.wiley.com/cgi-bin/abstract/112713559/ABSTRACT

Please contact:
Emanuel Lörtscher and Heike Riel, IBM Zurich Research Laboratory, Switzerland
E-mail: {eml,hei}@zurich.ibm.com

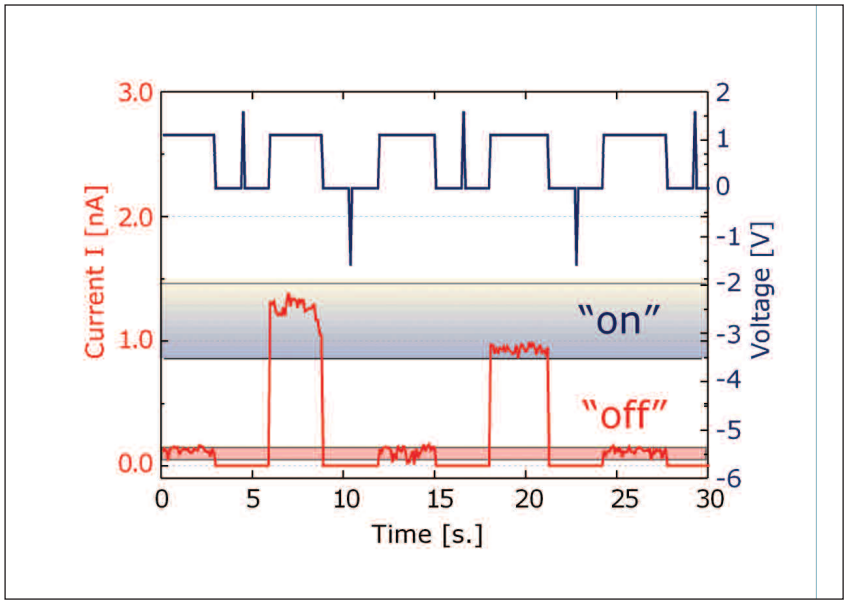

Figure 2:
Memory operation of the single-molecule system. The blue line shows the write, read, and erase pulse pattern applied and the red line demonstrates the resulting switching between "off" and "on" states of the molecular system. As can be seen in the picture, the system is initially in the "off" state. Then a write pulse of +1.6 V is applied and the molecule switches to the "on" state. This state can be read out using a voltage of +1.1 V. The molecule can be switched back to the "off" state by another pulse (erase pulse) of -1.6 V.

Mobile Hybrid Wireless Networks

HyMN: A Self-Organizing System for Interest-based Data Distribution in Mobile Hybrid Wireless Networks

by Steffen Rothkugel, Matthias R. Brust, Adrian Andronache

The concept of self-organisation already exists in the realm of ad-hoc networks. The inherent limitations of pure ad-hoc networks can be overcome by introducing uplinks to the Internet backbone. Extending the principles of self-organisation by transparently including those additional network links is the focus of the HyWercs project. HyMN, in turn, is a system aimed at fostering user experience by employing these concepts in the domain of multimedia content distribution.

Ad-hoc networks emphasise flexibility and survivability of the whole system. However, centralised approaches, for example for group management and information provisioning, do not work well in such settings. Moreover, due to frequent topology changes, connectivity of devices cannot be generally guaranteed. In particular, this makes it hard to disseminate information in a reliable way. We overcome these limitations inherent to pure ad-hoc networks by (a) establishing local groups of communicating devices in a self-organising manner and (b) introducing dedicated uplinks to a backbone infrastructure. Such uplinks are used for accessing resources available on the Internet. Additionally, they are employed to directly interconnect distant devices, either within a single partition and particularly also across different partitions. In practice, uplinks are realised for instance using cellular networks, satellites, or via Wi-Fi hotspots. Consequently, ad-hoc networks with devices that provide uplinks are called hybrid wireless networks throughout this paper.

One main goal of the HyWercs project is to develop a middleware that enables mobile devices to organise themselves in order to optimise the data access and provisioning process. The devices running the middleware keep track of the user's interests, which can be single data items or topics that comprise several data items. Devices sharing common interests will join interest groups. Whenever there is the need to retrieve information from a backbone network - like the Internet - or to establish a communication link to a device in another network partition, uplinks are employed. The mobile devices elect a particular member of the interest group, a so-called injection point, which maintains the uplinks.

As one proof of concept, the HyMN system (Hybrid Multimedia Network) has been implemented prototypically on top of the JANE1 simulation environment. JANE is designed to support application and protocol design in the realm of ad-hoc networks. One distinct feature of JANE is that applications can be executed on top of simulated devices (see figure) as well as in so-called platform mode on real hardware.

HyMN is designed for users interested in live multimedia news, for example from certain sports events. Mobile devices of football fans, for instance, create an interest group in a local ad-hoc network partition. Multiples of such groups coexist, such as football fans in pubs, those watching another match, traveling fans and more. In each of these cases, a considerable number of people have a shared interest and might join forces in a local setting.

The HyMN System in simulation mode.

The devices in the ad-hoc network running HyMN organise themselves in clusters, where the clusterheads act as injection points. For each group, interest in certain football matches is registered with the injection point. The injection point maintains an uplink to the backbone network in order to receive multimedia news related to the interests of the ad-hoc members. Thus, the football fans will receive injected information such as small videos, pictures or text messages each time something interesting is happening during the match. The multimedia files received remain stored on the mobile devices and will be provided to further devices joining the interest group later on.

In order to optimise the injection process, the backbone multimedia files are split into chunks. These chunks are injected concurrently to different injection points within a single ad-hoc network partition. Thus, the bandwidth for injecting data into a single partition is increased and the files reach the ad-hoc network faster. After reaching the network, the chunks are exchanged among the interested devices until all devices receive the complete file.

To realise the architecture presented, we focus on mobile devices establishing clusters in a self-organising way. For this, we have developed a clustering algorithm - namely WACA - that employs a heuristic weight function. The WACA algorithm (Weighted Application aware Clustering Algorithm) is designed to build an ad-hoc network topology that fits the needs of the application running on top of it. To achieve this, several parameters can be set in the weight function. The HyMN application focuses on multimedia content distribution from a backbone network to ad-hoc networks. To optimise this process, parameters like signal strength to the backbone network, long battery lifetime, dissemination degree and clustering coefficient are used for clusterhead election.

On-going research focuses on investigating how the mobile devices deal with different mobility scenarios. The HyMN system will be further developed and a middleware for wireless hybrid network devices will be developed from it.

Link:
http://mocca.uni.lu
http://syssoft.uni-trier.de/jane

Please contact:
Steffen Rothkugel,
University of Luxembourg
Tel: +352 466644 5259
E-mail: steffen.rothkugel@uni.lu

Visualization

Visualization of Colour Information on Highly Detailed 3D Models

by Matteo Dellepiane and Marco Callieri

It is now possible to display very detailed colour information on big 3D models. Research at the Visual Computing Lab, ISTI-CNR, Pisa, Italy, aims at creating a mostly automatic, easy-to-use way of mapping colour on geometric data.

The evolution of technology and important advances in the field of visualization of huge 3D datasets mean that it is now possible to acquire and display detailed 3D models. However, in order to achieve a completely realistic result, high quality colour information must be added to the geometric structure.

A coloured model can be created starting from: a 3D dataset, a set of photos (such as those shown in Figure 1) and calibration data, ie, the values of the parameters of the camera that took the photos. They can be divided into two groups: extrinsic parameters (translation vector and rotation matrix), related to the position of the camera in the space, and intrinsic parameters, the 'internal' settings (focal length, lens distortion) of the camera. Unfortunately, in most cases, calibration data are not available and have to be estimated.

Camera parameters are calculated by aligning each photo to the model: since this can only be done automatically in particular cases (eg, 'shape from silhouette'), user intervention is necessary.

Algorithms which align a photo on a model need some correspondences between them: the user must indicate some corresponding points on both objects.

Our software, called TexAlign, allows users to load both the model and all the photos, creating an Alignment Process whose data (correspondences coordinates, parameters of aligned images) are saved in an xml file. It is also possible to set correspondences not only between a photo and the model, but also between photos, using overlapping sections. The 'image to image' correspondence is used by the application to infer new correspondences with the model. A 'workload minimizer' analyses the graph of correspondences between the elements of the Alignment Process and helps the user to complete all alignments very quickly.

Once calibration data have been estimated, we have to 'project' the colour information onto the model. In 3D graphics, there are two ways to display colour on a model: texture mapping, where a single texture image is 'mapped' to the model, and 'per vertex' colour, where a RGB colour value is assigned to each vertex. The first technique gives very good results; the main limitation is that the photos have to be 'packed' in an image of maximum 4096x4096 resolution. Since we normally deal with tens of photos, we would need to subsample

Figure 1:
An example of a 3D model (centre image) with a set of photos. The result will be the 3D model with the colour information extracted from the photos.

them in order to put all of them in a single image. Hence, since we can provide very detailed 3D models, colour per vertex is the best choice to preserve both geometric and colour detail. We assign to each vertex of the model the colour of the images which project on it. Since more than one image could project onto the same vertex, a major issue rises: how can we automatically 'weigh' the contribution of each photo in order to achieve a realistic result?

To solve this problem, we developed an application called TexTailor. Starting from the set of photos, the model and the calibration data, TexTailor automatically creates a mask for each photo, assigning a weight to each pixel. The weight is calculated by considering three main values: the angle between the normal of the vertex associated with the pixel and the direction of view, the distance between the point of view and the vertex (depth) and the distance of the pixel from a discontinuity in the 'depth map' (that is a map where each pixel has the value of the depth of the associated vertex). These values are combined to calculate a weight. The colour value of each vertex is a 'weighted sum' of the contributions of all images. A result achieved using this approach is shown in Figure 2.

Even though the results so far are encouraging, more work is needed to improve the technique. The most promising area for future investigation is related to the estimation of the lighting environment. This can be done using known techniques (ie mirrored spheres) during the photographic campaign, or by combining image processing and analysis of geometric information during the projection operation. A good estimation of lighting would prevent the projection of dark areas in photos and could also provide information about the reflectance properties of the material.

Figure 2:
Details of the model of a wooden statue in Benabbio (Lucca): 36 photos projected on a 2.5 million face model.

Link:
Visual Computing Lab at ISTI-CNR:
http://vcg.isti.cnr.it/joomla/index.php

Please contact::
Matteo Dellepiane, Marco Callieri,
ISTI-CNR, Italy
Tel: +39 0503152925
E-mail: dellepiane|callieri@isti.cnr.it

Databases

Modular Access to Large Socio-Economic Databases - the Sedo Project in Luxembourg

by Uwe Warner, Thomas Tamisier and Fernand Feltz

Now online and regularly accessed, the Sedo homepage is the entry point to a mine of useful information about citizens' lives in Luxembourg. Underneath, the site is powered by a customisable platform for maintaining huge banks of data. Ergonomic search methods quickly deliver mains trends or detailed figures. Advanced users and people without statistical training can perform their own calculations with the guarantee of clear and accurate results.

The use of the Net for querying socio-economic databases has recently caught the attention of both experts and the public at large, with a number of online services provided by statistical organisations of various countries, such as those in Europe that belong to the CESSDA. Not only does the networked publication of statistics allow instantaneous update and correction, but it is also suited to building a relationship between the data provider and the data consumer. The consumer can deliver his feedback and formalise his requirements and the provider has the means to control the use of the data, so that the figures are better and more properly used. Other advantages are easing the retrieval of large data sets, allowing them to be customised, and enabling end-users to execute statistical functions for processing the data on the server.

The CEPS/INSTEAD is a public research institute for statistics and socio-economic studies dedicated to population, poverty and socio-economic policy. Through the Sedo (Socio-Economic Database Online) platform, set up in co-operation with the Centre de Recherche Public - Gabriel Lippmann, it offers a generalised and convenient access to the voluminous database on the national population. The information, collected through socio-economic household panels, consists of longitudinal data, which means that the same questions are periodically checked with the same respondents. Relevant domains are: demography, education, post-natal child-care, employment, equal opportunities, resources, housing, equipment and consumption and opinions about the monetary situation of the household and its evolution. Such exhaustive and disparate information is difficult to exploit due to the complexity and volume of the data. Up to now, only researchers with a solid background in longitudinal analysis could effectively benefit from this work, whereas it is relevant to all kinds of professional activities outside the area of the statistician experts. In addition, thanks to its precision and its constant updating, the socio-economic picture drawn by the combination of these miscellaneous themes is apt to significantly ease communication between the population and public servants and to support the decisions made by the political authorities.

Since autumn 2005, Sedo has been available online, featuring three different types of user access. A monitored free access for everybody to statistical tables and graphs on the themes of the study. A controlled access on demand using a user login and password that allows registered users to explore the banks of data without having to install statistical software on their own computers. An access of registered users, making a confidentiality pledge who are trained in statistical analysis to maintain individual and sophisticated statistics and research on the micro data. Finally, an Internet based discussion forum allows all users to exchange their experience and allows communication with those conducting the surveys and the research community. Compared to other information tools in the community, the Sedo project has great flexibility, a universal portability and the possibility of adapting to different databases. The most obvious benefit of the system is

Sedo's results for the public at large.

being able to obtain customised presentations of statistics with guaranteed accuracy and without the need for specialised statistical tools or an advanced expertise in statistical techniques.

For the sake of the reusability, it has been preferred to build an integrated platform rather than developing features on top of existing tools. This approach gives room for the integration of powerful statistical tools and lets the user perform online calculations according to the data. With every single variable or result are associated information types such as the statistical measurement (metric, ordinal, nominal), the population concerned (individuals, households), or the expected results for the variable. Based on these types, the system defines the statistical analysis that can be performed in order to ensure the coherence and soundness of the operations delivered to the user. To guarantee the coherence of the results and to ensure that they are correct with respect to the interpretation criteria of the user, Sedo also features an automatic calculation of the actual sampling size after a sequence of statistical operations.

From the programmer's point of view, full latitude is allowed for the integration of thesauruses and the customisation of the search-engine. The system is also characterised by a total freedom in the choice and the management of the databases.

With Sedo, the database server can be distinct from the server where the client requests are processed and the statistical operations are performed. Thanks to this independence of data storage and data processing, Sedo is portable to any kind of database. This separation is also useful for protecting the confidentiality of the data. In this case we must consider the different kind of users who will not be entitled to the same access to the data. It also minimises the volume of data used during the calculations, because only the data relevant to the requests is extracted and passed on to the processing phase.

Links:
http://sedo.ceps.lu
http://www.ceps.lu
http://www.lippmann.lu

Please contact:
Uwe Warner, Centre d'Etude de Population, de Pauvreté et de Politiques Socio-Economiques, Luxembourg
Tel: +352 58 58 55 554
E-mail: uwe.warner@ceps.lu

Thomas Tamisier, Centre de Recherche Public - Gabriel Lippmann, Luxembourg
Tel: +352 47 02 61 622.
E-mail: tamisier@lippmann.lu

Digital Libraries

BELIEF: Bringing Europe's Electronic Infrastructures to Expanding Frontiers

by Federico Biagini, Pasquale Pagano and Franco Zoppi

BELIEF aims at creating a platform where e-infrastructure providers and users can collaborate and exchange knowledge, thus bridging the gap separating the research infrastructure providers from the users. A key role is played by the BELIEF digital library, specifically designed to meet the needs of this community and to guarantee wide dissemination of its results.

The BELIEF digital library (DL) offers uniform access to multimedia documentation (eg technical reports, presentations, videos, manuals, on-line tutorials) providing continuously updated information on eInfrastructure-related projects, initiatives and events. The contents are harvested from different sources, such as projects web sites, repositories and databases. The DL provides services to support the submission, description, searching, browsing, retrieval, access, preservation and visualization of multimedia documents. Although designed to meet the needs of a specific community, the technology adopted by BELIEF can be easily adapted to meet the information and collaborative needs of other scientific communities.

In the BELIEF DL, users can define the information space they wish to search or browse in terms of the collections (ie, sets of documents) selected from those managed by the DL. Different search/browse options are offered: Google-like or fielded (with fields selected from a variety of known metadata formats), with or without relevance feedback. Users can search/browse any information associated with digital documents and their parts.

Supporting e-Infrastructure Projects

Projects can submit their documents directly to the digital library: BELIEF offers a Submission Interface and a Document Hosting Service to host projects' documents directly on the DL. To ensure that all documentation can be easily acquired, the DL stipulates a standard description format.

Projects using the Document Hosting Service do not need to bear any extra cost for the development of tools or services for their documents' submission, collection and retrieval. Projects that prefer to maintain their own documents can make them available to the community by simply providing a standard access to their documents and metadata so that they can be transparently 'harvested' by the DL.

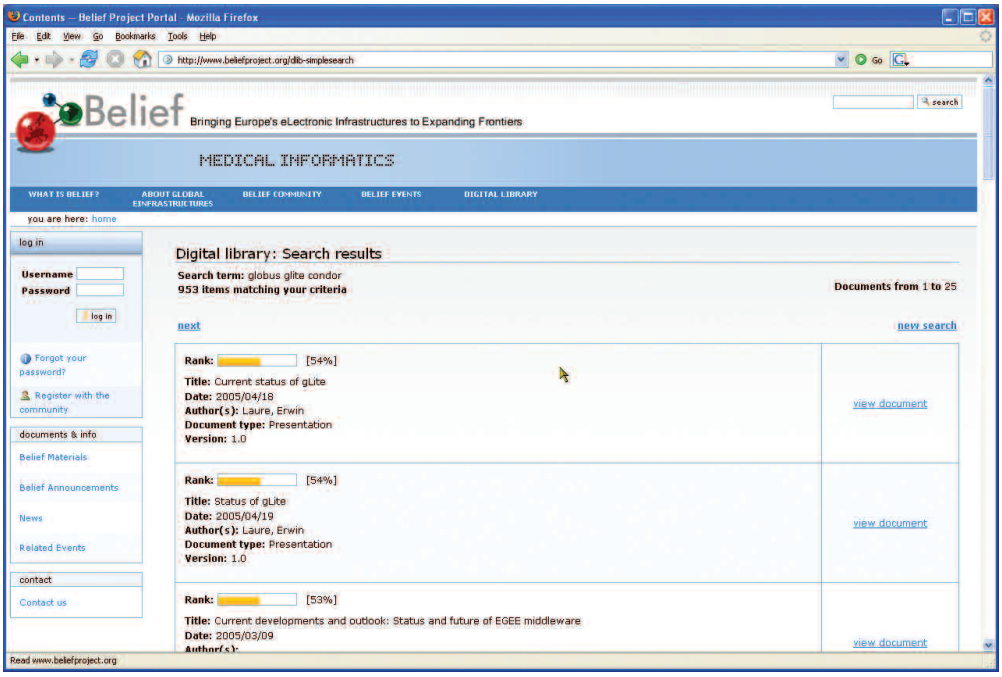

BELIEF digital library web interface.

Supporting Users' Community Needs

The BELIEF DL facilitates access to and dissemination of eInfrastructure project documents as it provides the user community with a "one-stop-shop" where all relevant documentation can be found. The DL provides a wide range of information on eInfrastructures, how they work, what projects exist and how to use them. Even documents stored in inner sections of the websites – and thus not indexed by general purpose search engines - are accessible via BELIEF. The user can search documents in a uniform way, and access documents produced by different providers via a unique interface. The DL knows the user's interest, as all users submit user profile information on initial registration. The DL can thus provide users with the most relevant documents according to their requirements and profile, thus providing an intelligent bridge between information provider and seeker. The figure shows the display of results for a search for documents containing information on the Globus and Glite systems plus the Condor project.

Technical Issues

We have concentrated on implementing harvesting tools that support the use of standard formats and conventions. BELIEF currently supports the following standard protocols and coding formats:

- Metadata encoding protocols: DC, DCQ (recommended), MARC, UNIMARC, MARC21.
- Metadata harvesting protocols: OAI-PMH - Open Archives Initiative Protocol for Metadata Harvesting (recommended) or any API call returning an XML file containing metadata encoded in one of the above mentioned encoding.
- File formats: XML (recommended), RSS.

Although information sources can be accessed via multiple protocols and formats, BELIEF has chosen to use Dublin Core Qualified encoding for metadata definition as it facilitates sharing of information between sources using different encoding systems with no loss of semantics.

A web-based interface allows easy submission of metadata and documents to the DL and 'document submission models' have been provided to cope with the specific characteristics of the documents managed by the DL (presentations, articles, books, tutorials, etc.). This model is easily customizable, even by end-users, and new models can be added by DL administrators at any time.

For developers interacting with the DL a set of APIs has been implemented. The BELIEF DL was created as an instance of OpenDLib (see Links), some of these APIs were already available as part of the OpenDLib Toolkit. However, in order to cope with emerging requirements a number of new APIs have been implemented; these are now included in the OpenDLib Toolkit (Browse, BrowseInfo, BrowseStandard, GetDocumentsPerTerm, Submit, ListCollections).

Future Work

Further activity is planned to add additional administration functionality to ensure smooth management of wider user communities nad more friendly help facilities for end-user. The Submission Service will be enhanced and the Access Service will be extended to support new protocols and formats to assist metadata compatibility over heterogeneous sources.

Links:

Networked Multimedia Information Systems Laboratory website:
http://www.isti.cnr.it/ResearchUnits/Labs/nmis-lab/
OpenDLib website:
http://www.opendlib.com
BELIEF project website:
http://www.beliefproject.org

Please contact:
Pasquale Pagano, ISTI-CNR, Italy
E-mail: pasquale.pagano@isti.cnr.it

Electromagnetic Radiation

Cell Phone Danger

by Martin Röösli and Harry Rudin

As we become more and more keen on cell phones there remains the open question of what effect these cell phones have on our health and well being. Some recent work in Switzerland sheds some light on these issues.

Our dependence on and investment in cell phones rise steadily. The annual sale of cell phones is approaching 1 billion units. Is the result of all the accompanying electromagnetic radiation going to have an effect on our health? Almost everyone has heard some 'anecdotal evidence' where there have been some undesirable effects on people: headache, concentration problems, irritability, insomnia, etc. But, can anything be said on a scientific basis? Some recent research from Switzerland has shed some light on the subject. More light is needed and some is coming.

What effect have cell phones on our health and well being?

The topic is difficult in itself and made more difficult by the enormous financial interests involved. Ideally, in a cold, scientific experiment, one should find a group of humans, some of whom would be subjected to the electromagnetic radiation produced by cell phones and their base stations for several years and some of whom in the group would not be exposed. All of them would be blind to exposure and they would not differ in terms of other factors influencing their health. One could carefully measure their health status and their well being, distinguishing between those exposed and those not exposed. Obviously this has not been done nor will it be done in the future.

While much material has been published with the objective of shedding light on the problem, many of the published studies have serious methodological limitations - leading to doubts concerning their validity. Reference [1] summarizes studies on mobile phone radiation and health-related quality of life, concluding, "The results are contradictory and the greater part of these studies is not able to address the issue of causality between exposure and outcome."

One of the first studies dealing with radio frequency electromagnetic field exposure was about the shutdown of the Swiss Schwarzenburg short-wave radio transmitter. While this took place back in 1998, a recently published paper [2] reports of measurements in the change in the melatonin produced by 54 volunteers living in the vicinity. The measurements showed that in the week after the shutdown, sleep quality improved and melatonin excretion increased by 15 percent, on average. (Melatonin plays an important role in regulating sleep-wake cycles because its production is inhibited by light and permitted by darkness. Melatonin is sometimes prescribed when suffering from jetlag.) However, the shut down date was known to the study participants. Thus, the authors state that "blinding of exposure was not possible in this observational study and this may have affected the outcome measurements in a direct or indirect (psychological) way."

In September, 2003, the Dutch TNO Physics and Electronic Laboratory announced the results of a study. They wrote: "The present study contributes to the research on finding a relation between electromagnetic fields and brain functions." Further, "Exposure to GSM-900 or GSM-1800 electromagnetic fields had no effect on well-being in either experimental group. However, exposure to a UMTS-like signal resulted in a small, but statistically significant impairment of well-being. Interestingly this effect was not only observed in individuals with self-reported health complaints attributed to daily life EMF exposures but also in a reference group without such complaints. This effect was found after only about half an hour's exposure to what, by everyday standards, was a relatively high environmental field strength (1 V/m). In practice, while individuals in the vicinity of operational UMTS antennas will be subject to continuous exposure, the field strengths in question will be lower." Thus, the transferability of the 'TNO'-results to normal day-to-day environmental exposure is questionable. [3]

There was considerable criticism of the circumstances of this study including that by its authors themselves. These authors and others recognized the need to make a more careful repeat of the

study. However, one result of the TNO study was a temporary and partial moratorium on the construction of new base stations in Switzerland, pending a more thorough study.

A Swiss scientific consortium was formed to carry out such a more thorough study. It was performed by a group of scientists from the University of Zurich, the University of Bern, and from the Swiss Federal Institute of Technology in Zurich. Several methodological improvements were implemented, including investigating a larger sample size and applying two different UMTS exposure levels, all yielding more robust results. The results were published on June 6, 2006. [4]

The effect of UMTS radiation found in the TNO study was not confirmed in the new Swiss study. It should be noted that the experiment had to do only with 45-minute UMTS-radiation exposure intervals (four of them separated by a period of one week) and radiation absorption in brain tissue was considerably smaller than during the use of a mobile phone. Thus, the authors concluded, "No conclusions can be drawn regarding short-term effects of cell phone exposure or the effects of long-term, base-station-like exposure on human health."[4] Nevertheless, several of the temporary and partial moratoria for the construction of new UMTS base stations in Switzerland were withdrawn, even though the question of long-term effects of UMTS radiation on humans remains open.

Meanwhile, the REFLEX study under Professor F. Adlkofer, was making measurements on isolated cells to see if these were effected by a high or low-frequency (power-line frequency) EMF (electromagnetic field). Quoting from the report: "The main goal of the REFLEX project was to investigate the effects of EMF on single cells in vitro at the molecular level below the energy density reflected by the present safety levels."[5] The study showed that in certain human cell types there was indeed a significant increase in the number of single- and double-strand breaks in the DNA as a result of high [6] and extremely low-frequency electromagnetic fields [5].

These results fascinated Professor Primo Schaer at the Center for Biomedicine at the University of Basel. In a talk at a meeting on June 29 of this year, organized by the Swiss Research Foundation for Mobile Communication, Professor Schaer gave a preliminary report on his own experiments which confirm the work in Vienna, showing that intermittent extremely low frequency fields can result in damage to DNA. At this point in time these research results are unpublished; but, when published, we will report on them here in ERCIM News. Professor Schaer emphasized that some DNA damages are repaired by the DNA repair mechanism. Thus, the observed genotoxic effects do not necessarily mean that EMF is carcinogenic for the human.

In conclusion, we now know that EMF does have some effects on humans and human cells. Since we still do not know what the effect is on our health and well-being, some degree of caution would seem to be called for.

References

**[1] Seitz H., Stinner D., Eikmann T., Herr C., Röösli M.; Electromagnetic Hypersensitivity (EHS) and Subjective Health Complaints Associated with Electromagnetic Fields of Mobile Phone Communication: A literature review [of papers] published between 2000 and 2004.
Science of the Total Environment, 2005; 349 (1-3): 45-55.**

[2] Altpeter, E.-S, Martin Röösli, M., Battaglia, M., Pfluger, D., Minder, C. and Abelin, T.; Effect of Short-Wave (6-22 MHz) Magnetic Fields on Sleep Quality and Melatonin Cycle in Humans: The Schwarzenburg Shut-Down Study, Bioelectromagnetics, Vo. 27, pp 142-150, 2006.

[3] Zwamborn, A.P.M., Vossen, S.H.J.A., van Leersum, B.J.A.M., Ouwens, M.A., Makel, W.N.; Effects of Global Communication System Radio-Frequency Fields on Well Being and Cognitive Functions of Human Subjects with and without Subjective Complaints. Netherlands Organization for Applied Scientific Research (TNO). FEL-03-C148, 2003.

[4] Regel S.J., Negovetic S., Röösli M., Berdiñas V., Schuderer J., Huss A., Lott U., Kuster N., Achermann P. ; UMTS Base Station Like Exposure, Well Being and Cognitive Performance. Environmental Health Perspectives 2006, 114 (8): 1270-1275 (http://www.ehponline.org/members/2006/8934/8934.html)

[5] Winker, R, Ivancsits, S., Pilger, A., Adlkofer, F. and Rudiger, H. W. Chromosomal Damage in Human Diploid Fibroblasts by Intermittent Exposure to Extremely Low-Frequency Electromagnetic Fields, Mutation Research, Vol. 585, Issues 1-2, pp. 43-49, 1 August 2005, http://www.verum-foundation.de/cgi-bin/content.cgi?id=euprojekte01

[6] Diem E., Schwarz C., Adlkofer F., Jahn O., Rüdiger H.; Non-Thermal DNA Breakage by Mobile Phone Radiation (1800 MHz) in Human Fibroblasts and in Transformed GFSH-R17 Rat Granulosa Cells in Bitro. Mutation Research, Vol. 583, (2), pp. 178-183, 2005.

Links:

TNO follow up study: http://www.mobile-research.ethz.ch/projekte.htm#18

EMF projects at the University of Bern: http://www.ispm.ch/index.php?id=814

Primos research group at the University of Basel: http://pages.unibas.ch/dbmw/biochemie/

Please contact:
Martin Röösli, University of Bern, Switzerland
E-mail: roeoesli@ispm.unibe.ch

Harry Rudin, Consultant, Switzerland
E-mail: hrudin@smile.ch

Satellite Communication Services

SAFE: Satellites for Epidemiology and Health Early Warning

by Catherine Chronaki and Laurent Braak

The risk of epidemics and emerging or re-emerging diseases such as avian flu or chikungunya fever is rising. Luckily, these risks can be contained with prevention, early warning, and prompt management. In this context, current early warning and response systems, although well-developed in Europe, could benefit from satellite services.

Detection and monitoring of potential risks has become an important part of the so-called 'Epidemic Intelligence' as a reaction of surveillance systems to the risk of epidemics arriving from less-developed areas outside Europe. Health early warning systems in areas remote, inaccessible, or prone to natural or man-made disasters enabled by satellite communications can significantly limit the risk of onset and the effects of epidemics and contribute to settling major public health issues. Economically speaking, satellite communications can also save money by enabling rapid and coordinated response and optimal adjustments of resources when deploying an emergency plan.

The SAFE project, co-funded by the European Space Agency (ESA), will develop and demonstrate the added value of satellite communication services, including low and high bandwidth access to Internet, co-operative working and geolocalisation for all phases of biological crisis including prevention, early warning, and crisis management. SAFE aims to establish a roadmap for ESA to determine how satellite services, by enabling or restoring access to information, can be integrated in European healthcare systems and civil protection authorities.

The SAFE project is part of the ESA telemedicine via satellite transition phase, set up to pave the way for the creation of a European user-driven telemedicine via satellite programme. ESA together with the World Health Organization (WHO) are joining forces for reaching this goal.

ESA through SAFE and with the assistance of WHO will specify, develop, and demonstrate a European outbreak early warning system that will become a key tool for the national and European bodies in charge of epidemiological surveillance and especially the European Center for Disease Control (ECDC). SAFE aims towards a better assessment of the epidemiological risks based on real-time objective data, which is a strategic objective not only for public health policy, but also for the eHealth industry. The SAFE solution that will be marketed to the different actors of epidemiological surveillance comprises the following components:

- *a network of expert centers and hospitals* able to communicate by satellite, and interoperable with other networks (eg gateways to Galileo and GMES-RESPOND);
- *a coordination van* equipped with a DVB-RCS capability and Internet access for communication with decision makers and expert centers. It has the role of the local coordination centre. Communication between the local coordinator and the mobile teams on the field will be ensured via a local Wi-Fi network and Satellite phones outside the Wi-Fi coverage. Mobile teams will be equipped with handheld terminals for data collection;
- *a field laboratory or lab kit* equipped with biological and biomedical equipment will identify microbial agents responsible for the threats to enable in-situ analysis;
- *a SAFE information system* including a data collection sub-system and a GIS. Besides geolocalisation, the GIS will enable visualizing the evolution of the epidemic and monitoring the alarm levels associated with the different geographical regions. This system will be based on relevant standards and will be interoperable with existing information systems, ie national health information system, deployable emergency response information system, epidemiological information systems and others. It will facilitate monitoring of indicators and other measures potentially useful in modeling the transmission patterns. Dissemination of this data can be done with the Internet capability via satellite.

An end-to-end case of early warning at the onset of an epidemic will be demon-

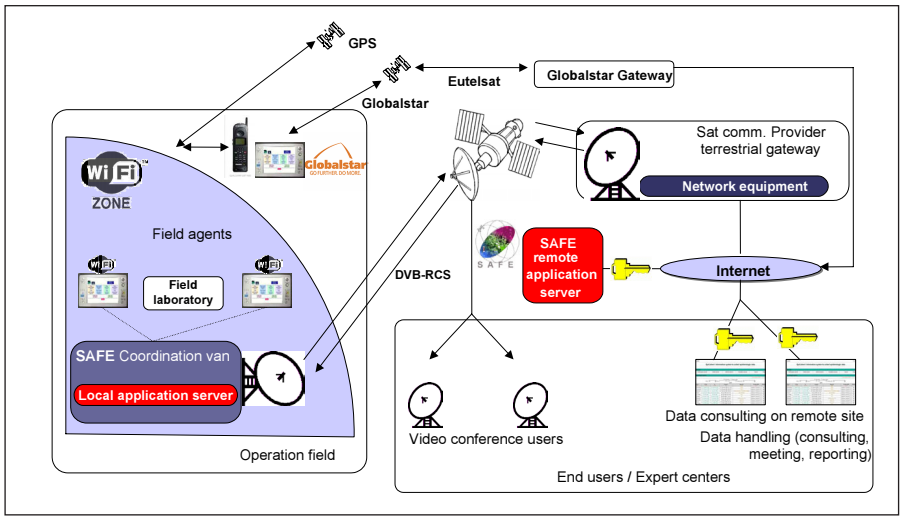

Components of the SAFE demonstrator.

strated at the Heraklion prefecture, in the island of Crete, Greece. The scenario based on epidemiological monitoring after an earthquake, namely the onset of typhoid fever epidemic, will validate the SAFE approach within the user community and if successful it could become part of regular earthquake readiness exercises. Furthermore, other demonstrations will be envisaged and implemented, namely related to avian flu and biological and radiological threats. These itinerant and on-demand demonstrations will be implemented all over Europe upon user request, promoting the added value of SAFE and allowing potential end-users and stakeholders to assess the demonstration according to relevant criteria, such as availability, interoperability, and security.

The SAFE consortium of nine partners from four countries is coordinated by MEDES (Institute for Space Medicine and Psychology, France). End users are represented by two epidemiology institutes - ISCIII (CNE) Instituto de Salud Carlos III 'Centro National de Epidemiologia' (Spain) and InVS 'Institut de Veille Sanitaire' (France) - to ensure compliance with policies and practices of epidemiological monitoring. ICS-FORTH (Greece) coordinates the demonstration of post-disaster biological risk monitoring in Crete. MICRO-UNIVr (University of Verona, Italy) is responsible for the mobile lab equipment and definitive identification of the microbial agent causing the biological crisis. MEDIVAN (Italy) provides industrial expertise on marketing mobile units for medical emergencies and rescue. GMV (Spain) ensures the overall system design, integration and the development of the GIS. TTSA (France) brings expertise in satellite communications and marketing of satellite services in health. REMIFOR (France) provides the coordination van and links with the professional bodies of civil protection and firemen. MEDES (France) coordinates the economical analysis and provides recommendations relying on its ongoing activities in epidemiology and humanitarian medicine involving satellite communications.

Links:

European Space Agency: http://www.esa.int

European Center for Disease Control: http://www.ecdc.europa.eu/

World Health Organization: http://www.who.int

Please contact:
Audrey Berthier,
MEDES, France (project coordinator):
E-mail: audrey.berthier@medes.fr

Pierluigi Mancini, ESA, Directorate of European Union and Industrial programs (Program manager)
E-mail Pierluigi.Mancini@esa.int

Catherine Chronaki, ICS-FORTH, Greece
E-mail: chronaki@ics.forth.gr

Satellite Devices

Planck Mission 70 GHz Receivers

by Jussi Varis

Planck is a European Space Agency - ESA mission to map the Cosmic Microwave Background (CMB) to an unprecedented sensitivity covering the whole sky. An accurate map of the CMB will allow the cosmologists to determine definitively the properties of our universe, like the age, the geometry, and the matter and energy composition. The Planck mission will be launched in 2008.

Finland has been responsible for building the crucial 70 GHz instruments for Planck. MilliLab, VTT has been leading the instrument development. Together with Elektrobit Microwave Ltd. and Metsähovi Radio Observatory, MilliLab delivered the six Planck low frequency instrument (LFI) 70 GHz front end modules (FEMs) and six back end modules (BEMs) for the mission (see figure) in 2006.

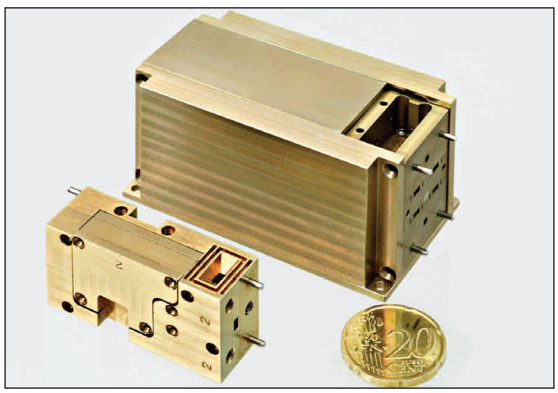

Planck LFI 70 GHz Front-End Module (FEM, in front)) and Back End Module (BEM, in back) developed by the Finnish Planck team. The FEM component houses four InP semiconductor low noise amplifiers and two phase shifters, and the BEM additional four low noise amplifiers.

The scientific performance of Planck is ultimately limited by the receiver noise level. To ensure very low noise performance, state-of-the-art Indium Phosphide (InP) semiconductors designed by the Finnish team were used in the 70 GHz FEMs and BEMs. The semiconductor devices were tested in MilliLab's unique cryogenic test system at the temperature of 20 K, which is the operating temperature of the FEMs in the Planck satellite. The FEMs and BEMs were manufactured and integrated in Elektrobit Microwave, and further tested at their cryogenic facility at the satellite operating conditions. The receivers built by the Finnish Planck team have demonstrated world record sensitivity at 70 GHz.

Please contact:
Jussi Varis, MilliLab-Millimetre Wave Laboratory of Finland
Tel: +358405694089
E-mail: jussi.varis@vtt.fi

Detector Technology

Data Acquisition in Scientific Applications

by Rob Halsall

Large scale scientific instrumentation systems for particle physics research at CERN are currently being deployed with many millions of sensor channels and Data Acquisition (DAQ) rates of the order of 100s of GByte/s. These systems, consisting of many racks of large format electronics modules, are typically implemented using custom digital logic in Field Programmable Gate Array (FPGA) devices. Installed in under ground caverns, closely coupled to massive detectors and networked to PC processing farms, these systems could be regarded as very large examples of embedded systems.

In other science areas (eg Synchrotron Radiation (SR), Neutron beams) there are growing requirements to move to similar detector technology. Although on a smaller scale this is resulting in high sample rates and channel counts with a dramatic increase in the data rates over previous systems from around 100 MByte/s to 10 GByte/s.

Recent advances in FPGA's have seen these devices evolve into programmable System On Chip (SOC) ASICs making these requirements feasible at reasonable cost and size. Typical FPGA components now contain multiple microprocessors, multiple multi-gigabit transceivers capable of driving common network standards, as well as multi million gates of programmable logic with DSP enhancements allowing performance rates up to 250 GMACS/s for the largest devices.

This provides the opportunity for shrinking our systems and producing embedded high performance DAQ with standard networked readout such as GBit Ethernet – Large Scale Science DAQ systems shrink wrapped into DAQ Sensor appliances.

In CCLRC Technology we have demonstrated the technical feasibility of embedded hardware based on high-speed serial optical links, standard network protocols, and programmable System On Chip (SOC) FPGA devices to deliver across a wide range of requirements and projects.

Our designs have implemented embedded hardware support features such as electronic fuses, supply monitoring, thermal monitoring, shutdown on fault, Wake On LAN type recovery from shutdown and low stand-by power. Remote reconfiguration of FPGA FLASH memory is possible together with encryption of the bit stream and automatic reversion to a fallback FLASH memory in the event of failure during re-programming. A typical small system might consist of a pixel sensor chip readout by an FPGA which is in turn readout by a PC server over a GBit Ethernet connection. A commodity SDRAM memory module connected to the FPGA provides data buffering and optical transceivers allow transmission of the GBit Ethernet over longer distances.

Figure 3 shows an embedded DAQ module planned for use on SR applications reading out 16 CMOS Pixel image sensors capable of a multi KHz frame rate. The module can log up to 800 MByte/s into SDRAM for a few seconds and then is readout to a PC server on two GBit Ethernet cables each sustaining 80 MByte/s.

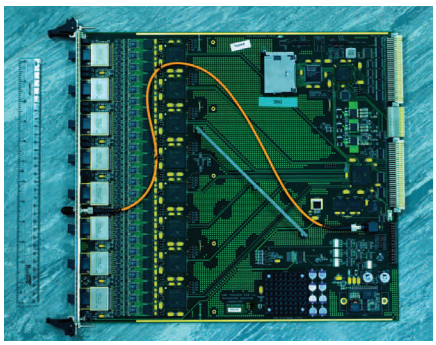

Figure 1:
One of 500 digitiser cards for a Large Scale Science Project.

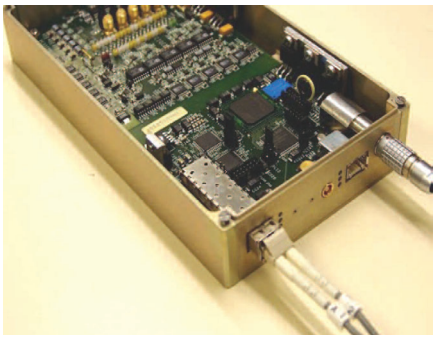

Figure 2:
High performance Embedded DAQ for Custom CMOS Sensor (not shown).

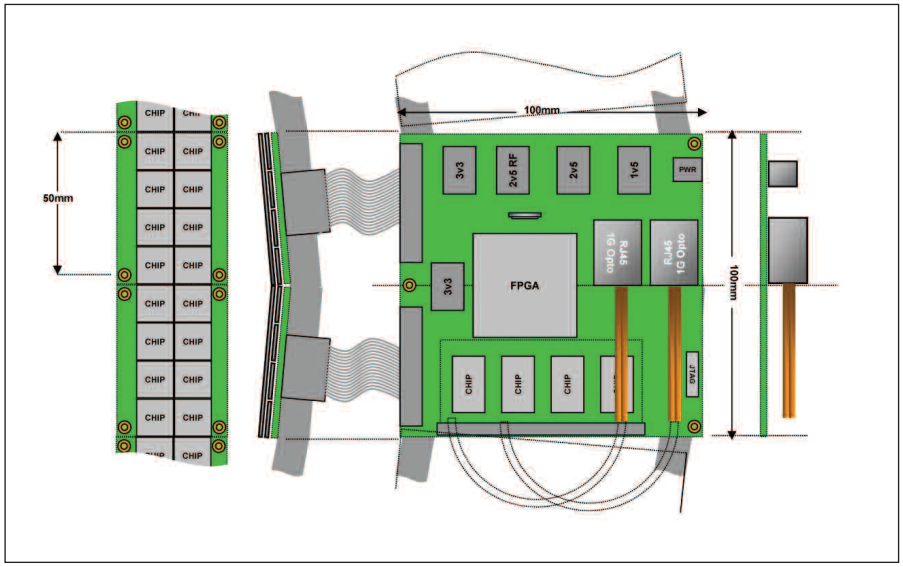

Figure 3: Multiple Custom CMOS Pixel Sensor with high performance Embedded DAQ.

One problem for these more widely deployed smaller projects based on high performance embedded modules is the design cost and turnaround time. Unfortunately there is no all-encompassing design tool for DAQ systems. The design process relies on an increasingly large and diverse portfolio of complex tools that require a high level of expertise to operate. Thus there is a growing cost of keeping up with the tool flows required to access more powerful hardware.

Currently we have an R&D program funded internally by our Centre for Instrumentation (CFI) which will investigate the embedded processor design tools. This project aims to reduce the cost of the firmware and software design for DAQ projects by developing a unified approach based around FPGA vendor embedded development tools.

The strategy we propose is to exploit the hardware IP modules and facilities existing within the embedded design tools by adding instrumentation-specific DAQ IP for integration of the sensor data stream. Customising the vendor's framework minimises the cost by reducing the amount of work that needs to be done, and minimises the risk by using standard interfaces.

A similar approach is taken for the software libraries - by using standard network protocols (Ethernet) and operating system APIs the software side of the DAQ library is more like 'DAQ middleware'.

The addition of embedded networked microprocessors within our designs gives us the ability to add higher levels of intelligence to our systems which might otherwise be difficult to include. Even allowing for the availability of C to hardware tools and an abundant gate count in the FPGA using embedded microprocessors would still be an efficient method.

Examples of the type of features which could be added include real-time tuning of the data flow and processing algorithms, logging of environmental information, safety critical monitoring of sensors (eg to prevent radiation damage), interaction with higher level control systems using standard protocols and provision of information to the e-science systems to aid in the management of the integrity of the data set.

Please contact:
Rob Halsall, CCLRC Technology
Rutherford Appleton Laboratory UK
Tel: +44 1235 445140
E-mail: R.Halsall@rl.ac.uk

Security

Network Anomaly Detection by Means of Machine Learning

by Roland Kwitt

Anomaly detection should be an integral part of every computer security system, since it is the only way to tackle the problem of identifying novel and modified attacks. Our work focuses on machine-learning approaches for anomaly detection and tries to deal with the problems that come along with it.

Since the number of reported security incidents caused by network attacks is dramatically increasing every year, corresponding attack detection systems have become a necessity in every company's network security system.

Generally, there are two approaches to tackle the detection problem. Most commercial intrusion detection systems employ some kind of signature-matching algorithms to detect malicious activities. In intrusion detection terminology this is called "misuse detection". Generally, such systems have very low false alarm rates and work very well in the event that the corresponding attack signatures are present. However, the last point leads directly to the potential drawback of misuse detection: missing signatures inevitably lead to undetected attacks!

Here, the second approach, termed anomaly detection, is in evidence. Anomaly detection maps normal behaviour to a baseline profile and tries to detect deviations. Thus, no a priori knowledge about attacks is needed any longer and the problem of missing attack signatures does not exist.

Our work deals with the applicability of machine-learning approaches, specifically Self-Organising Maps and neural networks in the field of network anomaly detection. By presenting a set of normal (problem-specific) feature vectors to a Self-Organising Map, it can learn the specific characteristics of these vectors and provide a distance measure of how well a new input vector fits into the class of normal data. In contrast to other machine-learning approaches, Self-Organizing Maps do not require a teacher who determines the favoured output.

Considering the aforementioned feature vectors, the vector elements strongly depend on the preferred layer of detection. That means, if anomaly detection is carried out at the connection level for instance, connection specific features, for example packet counts, connection duration, etc., have to be collected. However, if we require detection at a higher level, for example at application layer, another set of features will be necessary. Generally we can state, that the more discriminative power a feature possesses, the better detection results we get.

In case we use a neural network architecture for anomaly detection, such as a Feed-Forward Multi-Layer Perceptron (MLP), the training procedure for normal behaviour might not be apparent at first sight, since such networks gener-

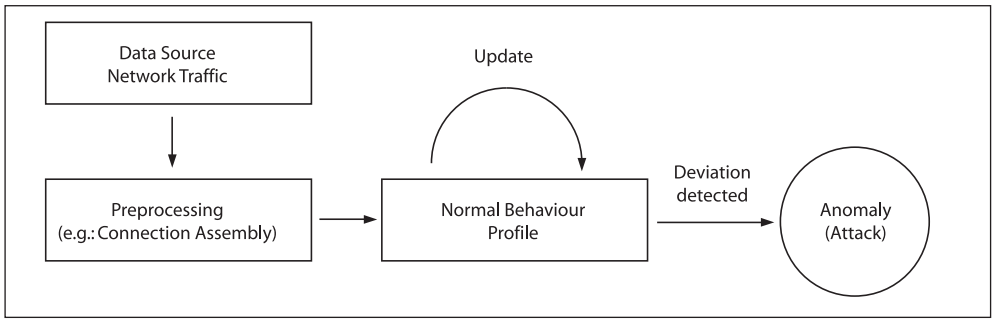

Simplified anomaly detection process.

ally require instances of both normal and abnormal data in the training phase. However, anomaly detection assumes that we only have instances of benign, normal data. This problem can be successfully solved by training the neural network with a set of random vectors, in the first step, pretending that all instances are anomalous. In the second step, the neural network is trained with the set of normal instances. In the ideal case, the net is then able to recognise recurrent benign behaviour as normal and classify potential attacks as abnormal. It is clear that we have to make the assumption that malicious behaviour is potentially abnormal.

One big problem that comes along with anomaly detectors is the lack of normal data in the training phase of a system. Most anomaly detection systems will miss an attack if the training data contains instances of that attack. However, no system administrator will waste his time, cleansing network traces from embedded attacks in order to produce adequate training data. What we need are detectors that are tolerant with regards to contaminated training data instances and still produce accurate results. Again, machine-learning with Self-Organising Maps can remedy that problem to a certain extent. Unless the amount of contamination is too high (which depends on the SOM parameters and the amount of normal data), no map regions will evolve which accept anomalous data.

Future work on this topic will include research on the feature construction process, which we found to be one of the most critical parts of an anomaly detector. Furthermore, we plan to evaluate machine-learning approaches that incorporate time as an important factor. Last but not least, we have to find a solution to the open problem of "concept drift", which denotes the fact, that normal behaviour changes over time, leading to high false positives (normal behaviour classified as anomalous) in most detection systems.

Please contact:
Roland Kwitt, Salzburg Research, Austria
Tel: +43 664 1266900
E-mail: rkwitt@salzburgresearch.at

Web 2.0 and Novel Interaction

From Solitary Podcasting to Ambient Pondcasting

by Emmanuel Frécon and Pär Hansson

Podcasting has taken the world by storm and become a mass phenomenon in less than a year. The core activities of podcast consumers, namely selecting broadcast sources and listening to portable media players, are solitary experiences. The Pondcaster project seeks to take podcasting a step further by encouraging group discovery and shared listening through a combination of interactive tabletop displays and personalised ambient listening stations.

Podcast was elected Word of the Year 2005 by the editors of the New Oxford American Dictionary and is now officially defined as "a digital recording of a radio broadcast or similar programme, made available on the internet for downloading to a personal audio player". In its current form, podcasting is a way for individuals and radio stations to broadcast shows by making them accessible as digital audio (and now also video) files over the internet. Listeners decide the type of programme they wish to hear and automatically have it downloaded so they can listen later, off-line, on portable media players such as the Apple iPod, which gave its name to the phenomenon.

Both the process of selecting shows to subscribe to for automatic synchronisation and that of listening at a later time are solitary experiences. While the various communication media offered by the internet can help in discovering and selecting shows, there are few ways for a co-located group to interactively share content and experience from a given set of already-loaded media players. This is a problem, especially as the venue of media players has been criticised as an invention that increases personal isolation in traditionally public or semi-public spaces such as streets, transport systems or even homes.

The SICS Pondcaster project addresses these problems and proposes a system that radically steps away from current podcast selection and consumption by supporting shared experiences. The Pondcaster system comprises a tabletop interactive artefact, which uses a pond environment as its metaphor, and any number of ambient listening stations placed in the environment to allow group listening, thus, bringing back the social

Figure 1: The Pondcaster interactive artefact and some mobile devices.

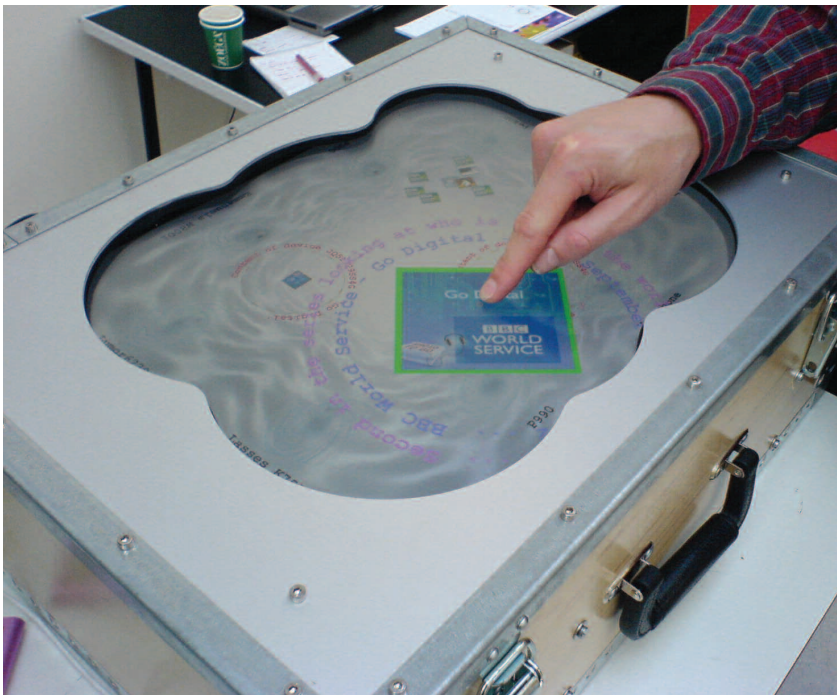

Figure 2: Upon interaction, creatures representing shows surface and the artefact starts playing their content.

aspects of radio listening from the early days of radio. The tabletop artefact builds on devices that can be found in every listener's pocket, such as media players and possibly mobile phones, and features unencumbered interaction via a touch screen.

The tabletop artefact acts as a sink for the podcast content of the users' media players, providing a snapshot of their current choices and podcast tastes. On physical connection, the content of any media player is automatically shown, using an animated aquatic metaphor. The pond shows groups of similar content as animated shoals of creatures, where each creature represents a podcast show. To avoid clutter, the metaphor is self-cleaning, whereby old, unused shoals slowly sink to the bottom of the virtual pond. By interacting with the creatures representing podcast shows, users can access show details and start listening, which favours collaboration and discovery of new sources. The artefact is also linked to a phone number and any user can send an SMS containing the words of a query. On reception, acting as an internet podcast search engine, the artefact will display matching podcast shows as a new shoal. These shows are typically not present on any connected media player, which again favours discovery of new shows. Finally, the artefact keeps track of bluetooth-enabled mobile phones in its close vicinity. Users can express their intention to listen further to newly discovered shows by dragging specific creatures to a representation of their phone.

Figure 3: Two users listening to podcast shows at one listening station. The TV screen shows the waiting list (some users off-picture) and sound comes from the directed speaker hanging from the roof.

The associations between shows and mobile phones will be remembered by the system. At a later time, users are able to listen to the chosen shows at public radio-like terminals. Users within the bluetooth vicinity of such a terminal will be recognised and their selections retrieved and played using a queue system. The setting encourages socialising, discussion and mobility since it keeps track of listening state across a set of listening stations for a given user.

The Pondcaster is based on existing internet technologies and because the system uses bluetooth tagging, no data is ever moved to mobile phones and media players. Subscriptions and preferences are kept on the Pondcaster and made available for further access by media players and Web browsers. As such, the artefact forms the premises of a personalised podcast aggregator.

Link:
http://www.sics.se/ice/projects/pondcaster

Please contact:
Emmanuel Frécon, SICS, Sweden
Tel: +46 8 6331500
E-mail: emmanuel@sics.se

EPoSS – A new European Technology Platform on Smart Systems Integration

by Erwin Schoitsch

The European Technology Platform on Smart Systems Integration (EPoSS) was launched in Brussels on 7th of July 2006. Information Society and Media Commissioner Viviane Reding officially opened the event and outlined in an encouraging manner the major issues to be tackled by EPoSS.

EPoSS brings together European private and public stakeholders in order to coordinate and to bundle efforts and to set-up sustainable structures for improving the competitiveness of European R&D on Smart Systems Technologies and their integration. Smart system integration is an emerging area with applications in many industrial sectors. A particular feature of the technology is its interdisciplinarity and the holistic viewpoint required for successful implementation. Systems combining sensing, processing and actuating functions are increasingly complex, involving a variety of disciplines and principles from physics, chemistry, engineering and biology. The systems may also integrate mechanical, electrical and biological functions in order to create new features like uninterrupted access to information, security and ease of use.

Photo by courtesy of EPoSS.

Information Society and Media Commissioner Viviane Reding.

New business opportunities - for example for automotive industries, a strong pillar of European economy (road safety, environment) - are also in the focus of EPoSS. Europe has a good position in micro/nanosystems, advanced technologies and embedded systems. Further progress, however, requires a change in the level and nature of these activities: regrouping the scattered industry and research communities and developing new instruments for interdisciplinary R&D activities as well as new models and formats of policy support for a fast transfer of research into innovative products.

Areas covered by EPoSS include:
- Automotive
- Aerospace
- Telecommunications
- MedTech Systems
- Logistics/RFID
- Commonalities (Cross Cutting Issues).

Rosalie Zobel, EC Director for Components and Systems welcomed the expected co-operation with the ETP ARTEMIS (embedded systems) and ENIAC (European Nanoelectronics Initiative Advisory Council).

Important industrial players are supporting and actively participating in EPoSS, such as EADS, Airbus, Thales, VW, Fiat, Siemens, Bosch, Infineon, to name just a few.

ERCIM members participating in EPoSS are Fraunhofer Gesellschaft and VTT. ARC Seibersdorf research (a member of the Austrian ERCIM member AARIT) will also join EPoSS.

The author, as the chairman of the ERCIM Working Group on Dependable Embedded Systems and member of the DECOS coordination team, was attending this event. Information on ERCIM and DECOS was distributed among the about 200 participants.

Link:
http://www.smart-systems-integration.org/

Please contact:
Erwin Schoitsch
ARC Seibersdorf research, Austria
E-mail: erwin.schoitsch@arcs.ac.at

ECDL 2006 - 10th European Conference on Digital Libraries

by M.Felisa Verdejo

The tenth European Conference on Digital Libraries was held at the University of Alicante, Spain from September 17 to 22, 2006. The event was jointly organised by the University of Alicante, Miguel de Cervantes Virtual Library and Universidad Nacional de Educación a Distancia (UNED).

The theme of this year's conference was Towards the European Digital Library. The aim was to emphasise the contribution made by the European Digital Library research community (and its liaisons with the international research community in general) to the current efforts of the European Commission in this direction. Speakers were invited to elaborate on this topic and the conference opened with a commemorative talk by Dr. Yannis Ioannidis looking back on the history of the ECDL. A CD-ROM is now available to the DL research community containing the proceedings of all the ECDL Conferences.

Although strongly international, the conference retained a distinctive European flavour. The 459 participants came from 45 countries: 66 percent were from Europe (26 countries), 23 percent from America, nine percent from Australasia and two percent from Africa.

The event followed the usual format of ECDL conferences consisting of a main conference, a doctoral consortium workshop, five tutorials, posters and demonstrations (this year there were 33), six workshops and the Cross-Language Evaluation Forum (CLEF), a major event on its own, with 140 participants.

The three guest speakers were Horst Foster, Director of Content in the Information Society and Media DG of the European Commission, who presented the Digital Libraries Initiative, one of the flagships of the i2010 strategic framework; Ricardo Baeza, director of Yahoo! Research Barcelona (Spain) and Yahoo! Research Latin America in Santiago (Chile), who focused on the potential of exploiting users' behaviour in search processes and Michael Keller, librarian at Ida M. Green University, Stanford, USA, who presented Google Book Search, a book-indexing project, and its benefits to readers and publishers, as well as a catalyst for other initiatives such as The European Digital Library.

A panel discussing the topic Sustained Digital Libraries for Universal Use was chaired by Ching-chih Chen of Simmons College, Boston, USA and José Borbinha of INESC-ID, Portugal. They were joined by Abdelaziz Abid from UNESCO, Vittore Casarosa from DELOS and Eric van der Meulen from the European Library.

Some 36 papers were selected for the main conference, an acceptance rate of 28 percent. Papers were organised in 12 sessions around the topics: Architecture (I, II), Preservation, Retrieval, Applications, Methodology, Metadata, Evaluation, User Studies, Modelling, Audiovisual content and Languages Technologies. The event was broadcast live on the Internet.

The doctoral consortium was scheduled the day before the conference at the same time as the tutorials in order to allow both seniors and students to fully participate in the conference itself. Two awards were made. The Best Paper Award, supported by the IEEE Technical committee on Digital Libraries (IEEE-TCDL) went to Carl Lagoze, Dean Krafft, Tim Cornwell, Dean Eckstrom, Susan Jesuroga and Chris Wilper for the paper Representing Contextualised Information in the NSDL, and the Best Young Researcher Award, supported by DELOS Network of Excellence on Digital Libraries, went to Daniel Coelho Gomes, Universidade de Lisboa, as coauthor of the paper Design and Selection Criteria for a National Web Archive.

As well as the scientific programme, ECDL 2006 included a full range of social activities allowing participants to

Ching-chih Chen of Simmons College, Boston, USA.

get together, and enjoy the local cultural heritage. On Sunday evening there was a boat trip around Alicante bay and the following evening there was a reception at the Santa Barbara fortress. The fortress lies on top of mount Benacantil and has splendid views over the harbour, the city and the coastline. Built by the Muslims in the ninth century, it was redesigned around 1580. On Tuesday evening there was a visit to the archaeological museum. This museum, which was European Museum of the Year in 2004, tells the history of the Mediterranean using installations, audiovisual projections and computer graphics. The conference closed with a dinner and jazz band. Delegates received a USB memory stick with photos of the conference events and further information about participants, sponsors and cooperating institutions.

This tenth conference highlighted research achievements in the field and provided a significant opportunity to demonstrate how the ECDL community is contributing towards the European Digital Library.

We are looking forward to the next ECDL Conference from September 16 to 21, 2007, in Budapest. For more information see http://www.ecdl2007.org/

The proceedings of ECDL 2006: Julio Gonzalo, Costantino Thanos, M.Felisa Verdejo , Rafael Carrasco (Eds). Lecture Notes in Computer Science, Vol 4172, 569 p.

The organisers of ECDL 2006 appreciate the financial support given to the Conference by Universidad de Alicante, Ministerio de Educación y Ciencia, Grupo Santander, Generalitat Valenciana, Patronato de turismo de Alicante and Universidad Nacional de Educación a Distancia (UNED).

Link:
http://www.ecdl2006.org

Cross-Language Evaluation Forum - CLEF 2006

by Carol Peters

The results of the seventh campaign of the Cross-Language Evaluation Forum were presented at a two-and-a-half day workshop held in Alicante, Spain, 20-22 September, immediately following the tenth European Conference on Digital Libraries. The workshop was attended by over 130 researchers and system developers from academia and industry.

The main objectives of the Cross-Language Evaluation Forum (CLEF) are to stimulate the development of mono- and multilingual information retrieval systems for European languages and to contribute to the building of a research community in the multidisciplinary area of multilingual information access (MLIA). These objectives are realised through the organisation of annual evaluation campaigns and workshops. The scope of CLEF has gradually expanded over the years. While in the early years, the main interest was in textual document retrieval, the focus is now diversified to include different kinds of text retrieval across languages and on different kinds of media (ie not just plain text but collections containing images and speech as well). In addition, attention is given to issues that regard system usability and user satisfaction with tasks to measure the effectiveness of interactive systems.

Evaluation Tracks

In CLEF 2006 eight tracks were offered to evaluate the performance of systems for:

- mono-, bi- and multilingual document retrieval on news collections (Ad-hoc)
- mono- and cross-language structured scientific data (Domain-specific)
- interactive cross-language retrieval (iCLEF)
- multiple language question answering (QA@CLEF)
- cross-language retrieval on image collections (ImageCLEF)
- cross-language speech retrieval (CL-SR)
- multilingual web retrieval (WebCLEF)
- cross-language geographic retrieval (GeoCLEF).

Test Suites

Most of the tracks adopt a corpus-based automatic scoring method for the assessment of system performance. The test collections consist of sets of statements representing information needs (queries) and collections of documents (corpora). System performance is evaluated by judging the documents retrieved in response to a topic with respect to their relevance (relevance assessments) and computing recall and precision measures.

The following document collections were used in CLEF 2006:

- CLEF multilingual comparable corpus of more than 2 million news documents in 12 European languages
- CLEF domain-specific corpora: English/German and Russian social science databases
- Malach collection of spontaneous speech in English and Czech, derived from the Shoah archives
- EuroGOV, ca 3.5 M webpages crawled from European governmental sites.

The ImageCLEF track used collections for both general photographic and medical image retrieval:

- IAPR TC-12 photo database; LTU photo collection for image annotation;
- ImageCLEFmed radiological database; IRMA collection for automatic image annotation.

Participation

Participation was up again this year with 90 groups submitting results for one or more of the different tracks: 60 from Europe, 14 from North America, 10 from Asia, 4 from South America and 2 from Australia.

Workshop

The campaign culminated in the workshop held in Alicante, 20-22 September. The workshop was divided between plenary track overviews, parallel, poster and breakout sessions. In her opening talk, Carol Peters, the CLEF Coordinator, stressed the need for more technical transfer activities. She commented that although many advances had been made in the multilingual information access research field there were still few real-world operational cross-language systems. In her opinion, CLEF should be paying more attention to issues that directly regard the user and the needs of the application communities rather than focusing most attention on system performance only in terms of precision and recall. In fact, one of the most interesting activities this year was the real-time question answering exercise, organised on-site by Fernando Llopis and Elisa Noguera, U. Alicante (see figure). Here the aim was to examine the ability of question answering systems to respond within a time constraint. The need for more technical transfer was taken up again in the final session in two talks. Martin Braschler, U. Applied Sciences Winterthur, Switzerland, gave an insightful talk on "What MLIA Applications can learn from Evaluation Campaigns" while Fredric Gey from U.C Berkeley, USA, summarised some of the main conclusions of the MLIA workshop at SIGIR 2006 in Seattle, where much of the discussion was concentrated on problems involved in building and marketing commercial MLIA systems. There was also an invited talk by Noriko Kando, National Institute of Informatics, Tokyo, Japan, on new evaluation activities at the NTCIR evaluation initiative for Asian languages.

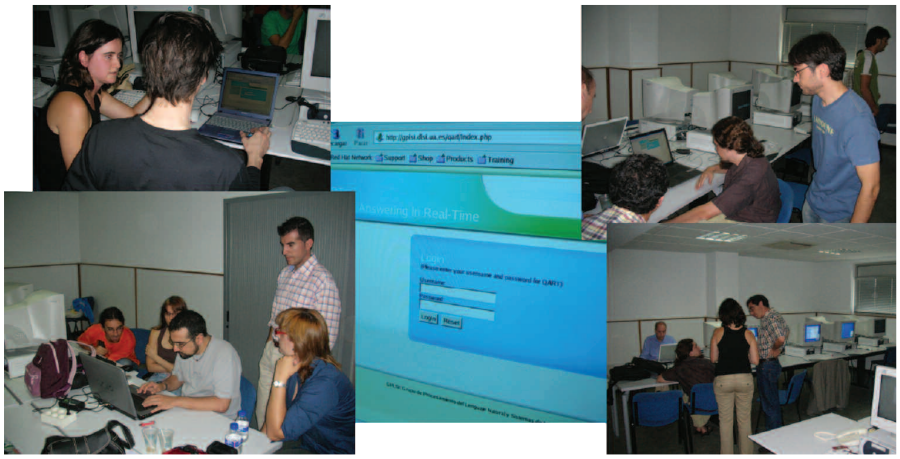

The real-time exercise: demonstrating the interface to the participants.

The presentations given at the CLEF Workshops and detailed reports on the experiments of CLEF 2006 and previous years can be found on the CLEF website at http://www.clef-campaign.org/. The preliminary agenda for CLEF 2007 will be available from mid-November. CLEF is an activity of the DELOS Network of Excellence for Digital Libraries.

Link:
http://www.clef-campaign.org

Please contact:
Carol Peters, ISTI-CNR, Italy, CLEF Coordinator
Tel: +39 050 3152897
E-mail: carol.peters@isti.cnr.it

ERCIM-DECOS Workshop

by Erwin Schoitsch

The ERCIM Working Group 'Dependable Embedded Systems' held a workshop in cooperation with the European Integrated Project DECOS - Dependable Embedded Components and Systems on 29 August 2006 at Euromicro 2006 in Cavtat, Dubrovnik, Croatia.

Embedded systems are regarded as the most important computer application field in the years to come, and a business driver for European industry. Therefore they are a targeted research area for European Research Programmes in Framework 7, promoted by ARTEMIS, the 'Technology Platform for Advanced Research and Technology for Embedded Intelligence'.

The workshop, entiteled 'Dependability Issues of Networked Embedded Systems: Research, Industrial Experience and Education', covered topics such as:

- Hardware/Software co-design, System-on-Chip integration
- Software processes for efficient development of dependable embedded systems
- Architectures, Methods and Tools for Design and Development of critical systems
- Interconnected, networked embedded systems
- Time triggered systems
- Functional Safety (system assessment, evaluation and standards)
- Verification, validation and certification of hardware and software, components and systems.
- Use/integration of COTS hardware and/or software as well as of legacy systems.

Attendance was free thanks to the sponsorship of the DECOS project.

Link:
ERCIM/DECOS Workshop:
http://www.itk.ntnu.no/misc/ercim/

Please contact:
Erwin Schoitsch, Austrian Research Centers/AARIT, Austria
E-mail: erwin.schoitsch@arcs.ac.at

ARTEMIS Annual Conference 2006

by Erwin Schoitsch

The third ARTEMIS annual conference took place in Graz, Austria on 22-24 May. The event was attended by nearly 250 delegates from European companies, representatives of SMEs, governments, academia and research Institutes.

ARTEMIS (Advanced Research and Technology for EMbedded Intelligence and Systems) is a strong, industry-driven European Technology Platform (ETP) whose goal is to establish a coherent, integrated European Strategic Research Agenda (SRA) to maintain the European lead in the field of embedded systems. From the scientific-technical point of view, ARTEMIS tries to establish a common approach and common technology across application domains. It does this by creating a public-private partnership bringing together people from industry, small and medium-sized businesses, universities, research centres and European public authorities, in the field of Embedded Systems. Common technology is defined by three working groups:

- *Reference Designs and Architectures:* to offer standard architectural approaches for major application contexts and to address complexity challenges, dependability issues and build synergies between application sectors (eg aeroplane safety at automotive costs).
- *Seamless Connectivity and Middleware:* to offer wide-scale interoperability and connectivity, to provide new services, functionality and to build the ambient intelligent (smart) environment
- *System Design Methods and Tools:* to enable dependable, rapid design and development.

Additionally, a strategic group works on generic enabling technologies based on foundational sciences' results.

Application contexts are wide-spread, according to the principle that Embedded Systems are Everywhere.

The target application areas are:

- Industrial (automotive, aerospace, railways, machinery, process industry, biomedical, manufacturing, medical systems, etc.)
- Nomadic Environments (enabling devices such as PDAs, on-body systems) for the mobile environment
- Private Spaces (homes, cars, offices, entertainment, comfort, well-being, safety)
- Public Infrastructures (airports, highways, railways, communication networks, security systems, energy, etc.).

ERCIM at the Smart Systems/ TTTech Booth at the ARTEMIS Conference Exhibition, Graz 2006.

ARC Seibersdorf research was represented at the exhibition by a booth displaying animated videos on the application of dependable embedded systems at NASA (time-triggered systems for the Mars mission), DARPA Grand Challenge (Embedded Vision Systems being part of one of the finalists in this race of autonomous vehicles through the desert in Nevada) and the DECOS project. The EU-funded Integrated Project DECOS - Dependable Embedded Components and Systems - was presented via videos explaining the DECOS technology in the context of automotive applications and by a demonstrator TT-Car, a model car demonstrating the advantages of time-triggered technology.

ERCIM and AARIT, the Austrian member of ERCIM, was represented at the DECOS project booth. The ERCIM Working Group, Dependable Embedded Systems (DES-WG), is cooperating closely with DECOS. The DES-WG is a dissemination partner of DECOS and DES-WG events are sponsored by DECOS.

Links:

ARTEMIS: http://www.artemis-office.org/

DECOS: https://www.decos.at/

Please contact:

Erwin Schoitsch, ARC Seibersdorf research/ AARIT, Austria

E-mail: erwin.schoitsch@arcs.ac.at

XP2006 - Extreme Programming and Agile Processes in Software Engineering

by Outi Salo and Pekka Abrahamsson

The 7th International Conference on eXtreme Programming and Agile Processes in Software Engineering was organized in Oulu, Finland in June 2006 by VTT Technical Research Centre of Finland and the University of Oulu.

As the largest agile conference ever organized in Europe, the five day program included altogether 18 tutorials, five keynote speeches, 22 research and experience paper presentations, 20 different workshops and activities, three cutting edge panel discussions and poster demonstrations, as well as 20 round-the-clock open-space events self-organized by the enthusiastic agile crowd. All together, 283 industrial and academic professionals from 25 countries around the world, and from over 100 different companies, gathered together to discuss their needs and ideas for incorporating agile methodologies into the production models. In this respect, the conference can undoubtedly be regarded as a huge success in terms of both quality and quantity!

Barry Boehm (left) and Kent Beck at XP2006.

The conference was also honoured with the presence of two legends in software industry; Barry Boehm (Professor of Software Engineering, Computer Science Department Director at USC Center for Software Engineering and creator of, eg, COCOMO, Spiral Model, and the Theory W approach), and Kent Beck (Founder and Director of Three Rivers Institute, the father of Extreme Programming methodology and one of the founders of the Agile Manifesto).

In all, the success of XP2006 is one indicator that agile methodologies and their variants are here to stay. The enormous and ongoing interest of both the industry and academia shows no signs of subsiding. However, it is also evident that even though the agile ideologies and methodologies have emerged since the mid 1990's, there is still a great need to keep the debate and evolution moving upwards and onwards. In fact, a majority of software organizations functioning in complex human and technical ecosystems are still contemplating whether to make the transition towards agile software development and how they should actually proceed in pursuit of increased quality and customer satisfaction, and decreased lead-time and costs.

XP2007 will be organized in June, 2007 in Como, Italy. Welcome to join the agile crowd!

Links:

http://virtual.vtt.fi/inf/pdf/symposiums/2006/S241.pdf (Collection of Tutorials, Workshops, Activities and Key Note Speeches of XP2006)

http://www.xp2006.org

http://agile.vtt.fi

Please contact:

Outi Salo, Pekka Abrahamsson

VTT Technical Research Center of Finland,

E-mail: outi.salo@vtt.fi,

pekka.abrahamsson@vtt.fi

Call for Participation

GRIDS@work: CoreGRID Conference, Grid Plugtests and Contest

Sophia-Antipolis, France, 27 November - 1st December 2006

Following the successful second GRID Plugtests in 2005 with 230 participants from many different countries, ETSI, INRIA and CoreGRID decided to organize this year an enhanced GRIDs@work event, composed of:

- a series of conferences and tutorials: view the programme
- a 3rd GRID Plugtests, during which two competitions will take place: the N-Queens Contest and the FlowShop Contest.

This event will take place on a Grid infrastructure. You may provide machines to this effect.

More information:
http://www.etsi.org/plugtests/Upcoming/GRID2006/GRID2006.htm

Call for Participation

SOFSEM '07 - Current Trends in Theory and Practice of Computer Science

Harrachov, Czech Republic, 20-26 January 2007

SOFSEM (originally SOFtware SEMinar) is the annual, international conference devoted to the theory and practice of computer science. Its aim is to foster cooperation among professionals from academia and industry working in all modern areas of computer science. The SOFSEM program consists of series of invited talks, given by prominent researchers and professionals, Contributed Talks selected by the Program Committee from submitted papers, Working Sessions discussing work-in-progress, and the Student Research Forum.

SOFSEM offers a unique opportunity to quickly obtain a representative and in-depth overview of the areas of computer science which were selected as the topics of this year. For an optimal atmosphere of close professional interaction, SOFSEM is organised in tracks with well-defined topics so as to attract all researchers and professionals that share specific fields of interest. SOFSEM is especially suited for young computer scientists.

Tracks

SOFSEM 2007 tracks are:

- Foundations of Computer Science
- Multi-Agent Systems
- Emerging Web Technologies
- Dependable Software and Systems.

As it has became a tradition in the last years, the proceedings with the invited and contributed papers will be published in the series Lecture Notes in Computer Science (LNCS, Springer-Verlag) which will be distributed at the conference.

SOFSEM 2007 is sponsored by ERCIM.

More information:
http://www.cs.cas.cz/sofsem/07/

Call for Papers

CAISE '07 - The 19th International Conference on Advanced Information Systems Engineering

Trondheim, Norway, 11-15 June 2007

CAiSE'07 invites submissions on the development, maintenance, procurement and usage of information systems. This year's special theme is 'Ubiquitous Information Systems Engineering', but as usual the conference also welcomes papers on other aspects of IS engineering. New information systems span different organizations and geographical locations, and often need to support the untethered mobility of the people using the systems. They have a large impact on the every day life of organisations as well as individuals. In the light of the development towards ambient, pervasive and ubiquitous computing this impact will increase significantly.

Topics

In addition to the special theme, relevant topic include:

- Methodologies and approaches for IS engineering:
- Innovative platforms and architectures for IS engineering
- Engineering of specific kinds of IS
- Quality concerns in IS engineering.

Important Dates

- 30 November 2006: Paper submission deadline
- 10 February 2007: Notification of acceptance
- 11-15 June 2007: Conference and workshops

The proceedings will be published in the series Lecture Notes in Computer Science (LNCS, Springer-Verlag).

CAISE 2007 is organised by NTNU and sponsored by ERCIM.

More information:
http://caise07.idi.ntnu.no/

Call for Papers

ECOOP 2007 - 21st European Conference on Object-Oriented Programming

Berlin, 30 July- 3 August 2007

The ECOOP 2007 conference invites high quality papers presenting research results or experience in all areas relevant to object technology, including work that takes inspiration from or builds connections to areas not commonly considered object-oriented.

Topics

ECOOP wishes to embrace a broad range of topics, so the following list of suggested topics is by no means exclusive:

- Architecture, Design Patterns
- Aspects, Components, Modularity, Separation of Concerns
- Collaboration, Workflow
- Concurrency, Real-time, Embeddedness, Mobility, Distribution
- Databases, Persistence, Transactions
- Domain Specific Languages, Language Workbenches
- Dynamicity, Adaptability, Reflection
- Frameworks, Product Lines, Generative Programming
- HCI, User Interfaces
- Language Design, Language Constructs, Static Analysis
- Language Implementation, Virtual Machines, Partial Evaluation
- Methodology, Process, Practices, Metrics
- Model Engineering, Design Languages, Transformations
- Requirements Analysis, Business Modeling
- Software Evolution, Versioning
- Theoretical Foundations, Formal methods
- Tools, Programming environments.

Research papers should advance the current state of the art. Many different research methods can be applied, eg, both experimentally based work and mathematical results are welcome. Experience papers should describe novel insight gained from the practical application of object technology, in such a way that it is of interest to a broad group of researchers and practitioners.

Workshops

ECOOP 2007 will host a number of workshops addressing different areas of object-oriented technology. Workshops serve as a forum for exchanging late breaking ideas and theories in an evolutionary stage.

Important Dates

- Submission of papers: 13 December 2006
- Update period: 13-20 December 2006
- Deadline for Proposals: 20 December 2006

ECOOP 2007 is sponsored by ERCIM.

More information:
ttp://ecoop07.swt.cs.tu-berlin.de/

Call for Papers

The Sixth Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering

Dubrovnik, Croatia, 3-7 September 2007

The biennial series of joint ESEC and FSE provides a lively and outstanding forum where researchers and practitioners can report and discuss recent research results and trends, as well as their impact on practical application in all areas of software engineering. ESEC/FSE yields the just-right scope, wide enough to portray a comprehensive and truly-international state of the art overview, but small enough to keep an interactive and productive atmosphere.

ESEC/FSE has continuously attracted participants from industry and academia, confirming that the venue provides topics that are up-to-date, significant, and interesting. ESEC/FSE 2007 intends to continue this tradition by offering a strong technical program associated with an exciting list of events, including keynote talks by leaders in the field, invited state-of-the-art seminars, workshops, a doctoral symposium, and the new Widened Software Engineering (WISE) presentations.

Topics

Contributions to the above events on any field of Software Engineering are invited, including, but not limited to:

- Requirements Engineering
- Specification and Verification
- Software Architecture and Design
- Analysis and Testing
- Development Paradigms and Software Processes
- Evolution and Refactoring
- Tools and Environments
- Empirical Software Engineering
- Software Metrics
- Software Quality and Performance
- Component-based Software Engineering
- Model Driven Engineering
- Web Applications
- Distributed Systems and Middleware
- Service Oriented Applications
- Mobile and Embedded System
- Open Standards and Certification
- Software Economics and Human Resources
- Dependability (safety, security, reliability)
- Case Studies and Experience Reports.

Important Dates

- Workshop proposals submission: 5 March 2007
- Workshop proposals notification: 15 March 2007
- Technical papers submission: 19 March 2007
- WISE papers submission: 3 April 2007
- Doctoral Symposium abstracts submissions: 14 May 2007
- WISE papers notification: 3 May 2007
- Technical papers notification: 1 June 2007
- Doctoral symposium notification: 10 June 2007
- Camera-ready contributions (all): 1 July 2007

The proceedings will be published in the series Lecture Notes in Computer Science (LNCS, Springer-Verlag).

ESEC/FSE 2007 is sponsored by ERCIM.

More information:
http://www.idt.mdh.se/esec-fse-2007/

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 10,500 copies. The printed version of ERCIM News has a production cost of 8 Euro per copy. It is available free of charge for certain groups.

Advertising
For current advertising rates and conditions, see http://www.ercim.org/publication/ERCIM_News/ or contact office@ercim.org

ERCIM News online edition
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: Jérôme Chailloux, ERCIM Manager

Central Editor:
Peter Kunz, ERCIM office
peter.kunz@ercim.org

Local Editors:
AARIT: n.a.
CCLRC: Martin Prime
M.J.Prime@rl.ac.uk
CRCIM: Michal Haindl
haindl@utia.cas.cz
CWI: Annette Kik
Annette.Kik@cwi.nl
CNR: Carol Peters
carol.peters@isti.cnr.it
FORTH: Eleni Orphanoudakis
eleni@ics.forth.gr
Fraunhofer ICT Group:
Michael Krapp
michael.krapp@scai.fraunhofer.de
FNR: Patrik Hitzelberger
hitzelbe@lippmann.lu
FWO/FNRS: Benoît Michel
michel@tele.ucl.ac.be
INRIA: Bernard Hidoine
bernard.hidoine@inria.fr
Irish Universities Association:
Ray Walshe
ray@computing.dcu.ie
NTNU: Truls Gjestland
truls.gjestland@ime.ntnu.no
SARIT: Harry Rudin
hrudin@smile.ch
SICS: Kersti Hedman
kersti@sics.se
SpaRCIM: Salvador Lucas
slucas@dsic.upv.es
SZTAKI: Erzsébet Csuhaj-Varjú
csuhaj@sztaki.hu
VTT: Pia-Maria Linden-Linna
pia-maria.linden-linna@vtt.fi
W3C: Marie-Claire Forgue
mcf@w3.org

Subscription
Subscribe to ERCIM News by:
- contacting the ERCIM office (see address above)
- filling out the form at the ERCIM website at http://www.ercim.org/

News about legal information from European directives, and pan-European legal requirements and regulations.

Commission released the current draft of Proposal for a Council Framework

Decision on the protection of personal data processed in the framework of police and judicial co-operation in criminal matters to the Secretary-General of the Council. The current draft says that data can only be exchanged with non-EU states "if an adequate level of data protection is ensured in the third country or by the international body to which the data concerned shall be transferred". Five member states support this 'adequacy' test (Czech Republic Switzerland, Finland, Greece and Portugal) but seven - backing the US complaint - are opposed to its inclusion (Germany, Denmark, Spain, Ireland, Norway, Sweden and UK) 24 Aug 2006 the original draft see, http://europa.eu.int/rapid/pressReleasesAction.do?reference=MEMO/05/349&format=HTML&aged=0&language=EN&guiLanguage=en

EU: Visa Information System (VIS) takes shape

The EU is currently deciding on the purpose, function and scope of the VIS and law enforcement access to it. The personal data of everyone who applies for an EU short-stay visa, including their photograph and fingerprints, will be recorded in the VIS (this includes persons whose applications for a visa are rejected). As with the "US VISIT system", this data will ultimately be used to facilitate identity checks and verify entry to and exit from the EU. In addition, the "internal security" agencies of the member states and Europol will have access to the data. See Draft Regulation on VIS and the exchange of data between Member States on short-stay visas, 7 September 2006, See also http://www.libertysecurity.org/article108.html

Privacy International released a comprehensive Global Freedom of Information Survey.

Privacy International provides a comprehensive review of Freedom of Information Laws and practices in nearly 70 countries around the world. the survey draws attention to the growing movement around the world to adopt FOI laws. While it also highlights that many problems still exist such as poorly drafted laws, lax implementation and an ongoing culture of secrecy in many countries. 28 September 2006 the survey is available at http://www.privacyinternational.org/foi/survey

The European Parliament adopted 'European Parliament resolution on the interception of bank transfer data from the SWIFT system by the US secret services'.

The European Parliament, demanding explanations from EU governments and institutions regarding their complicity in the SWIFT transfers of financial data to the US on 5 July 2006. See, http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//TEXT+TA+P6-TA-2006-0317+0+DOC+XML+V0//EN

edited by Yue Liu, NRCCL, University of Oslo, Norway

ISTI-CNR - **Donatella Castelli coordinated the first Digital Library Reference Model workshop** in Frascati, near Rome, in June. At the workshop, an invited group of international experts discussed the digital library (DL) model proposed by DELOS, the Network of Excellence on Digital Libraries. Despite the large number of digital library systems in existence, there is so far no clear agreement on what these systems are exactly, nor what basic functionality they should provide. This model is a first attempt at establishing an abstract DL framework, based on a small number of unifying concepts. It seeks to provide a common semantics that can be used across different DL implementations. The workshop aimed at achieving global consensus on this model. In her introductory speech, Patricia Manson, Head of the EC unit for 'Technology-Enhanced Learning & Cultural Heritage' said that she expected that the results of the workshop would impact on the definition of FP7 research priorities and help to initiate the relevant political and operational actions needed to deliver the European Digital Library to European citizens.

Donatella Castelli.

CWI - **Christine Bachoc, Université Bordeaux, and Frank Vallentin, CWI, have found new upper bounds for 'kissing' in higher dimensions**. In geometry the kissing number is the maximum number of non-overlapping unit balls that can simultaneously touch a central unit ball. In two dimensions the kissing number is six. This can be seen when euro coins are grouped around one central euro coin (see picture). The kissing number is only known for the dimensions 1, 2, 3, 4, 8 and 24. For the dimensions 5, 6, 7, 9 and 10 Bachoc and Vallentin now found sharper upper bounds. The kissing problem has a rich history. In 1694 Isaac Newton and David Gregory had a famous discussion about the kissing number in three dimensions. Gregory thought thirteen balls could fit while Newton believed the limit was twelve. Only in 1953, Schütte and Van der Waerden proved Newton right.

Kissing coins.

An article explaining the results is available at http://arxiv.org/abs/math.MG/0608426.

Fraunhofer-Gesellschaft - **Ulrich Trottenberg** took over the presidency of the 'Fraunhofer Center of Institutes', Schloss Birlinghoven, on 1st July 2006. The 'center of institutes Schloss Birlinghoven' gathers all research institutes at the Birlinghoven campus. With some 500 scientists, it is the largest center for information technology of the Fraunhofer Gesellschaft. During his presidency Trottenberg intends to make the leading role of the institute center more visible. He is planning to do this by numerous events/meetings for industry, research, politics and the general public.

Ulrich Trottenberg.

Ulrich Trottenberg holds the chair for Applied Mathematics and Scientific Computing at the Department for Mathematics at the University of Cologne. He also is the head of the Fraunhofer Institute for Algorithms and Scientific Computing (SCAI) and vice chairman of the Fraunhofer Research Group for Information and Communication Technology (ICT). In addition he heads the Department of Simulation and Software Technology of the German Aerospace Center.

Current research topics of the Fraunhofer Center of Institutes are the extension of the world-wide web to information technology services, the handling of the data flood with automatic extraction of knowledge from data (datamining) and Human-Machine Interaction.

A team of researchers from the Politecnico di Milano and CEFREIL Labs, led by **Stefano Ceri**, won the 2006 Semantic Web Services Challenge (SWS), organised by Stanford University. The goal of the SWS Challenge is to develop a common understanding of various technologies intended to facilitate the automation of mediation, choreography and discovery for Web Services using semantic annotations. The challenge required the solution of a Web Service mediation problem and a Web Service discovery problem. The implementation presented by the Italian team was judged the most complete.
See:
http://deri.stanford.edu/challenge/2006/

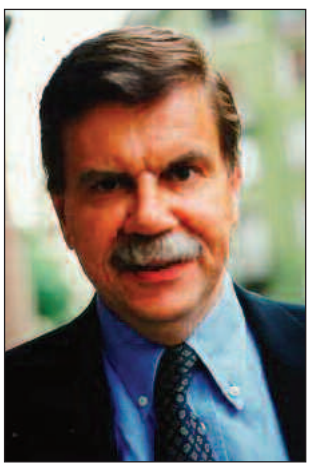

Stefano Ceri.

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/

Council for the 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.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
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
Tel. +352 26 19 25-1, Fax +352 26 1925 35
http:// www.fnr.lu

FWO
Egmontstraat 5
B-1000 Brussels, Belgium
Tel: +32 2 512.9110
http://www.fwo.be/

FNRS
rue d'Egmont 5
B-1000 Brussels, Belgium
Tel: +32 2 504 92 11
http://www.fnrs.be/

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/

Spanish Research Consortium for Informatics
and Mathematics c/o Esperanza Marcos, Rey Juan Carlos
University, C/ Tulipan s/n, 28933-Móstoles, Madrid, Spain,
Tel: +34 91 664 74 91, Fax: 34 91 664 74 90
http://www.sparcim.org

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/

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

Irish Universities Association
c/o School of Computing, Dublin City University
Glasnevin, Dublin 9, Ireland
Tel: +3531 7005636, Fax: +3531 7005442
http://ercim.computing.dcu.ie/

Technical Research Centre of Finland
PO Box 1000
FIN-02044 VTT, Finland
Tel:+358 207226041, Fax :+207226027
http://www.vtt.fi/

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
2004 route des Lucioles
BP 93
F-06902 Sophia Antipolis Cedex
Fax: +33 4 9238 5011
E-mail: office@ercim.org

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

I wish to subscribe to the

☐ ***printed edtion*** ☐ ***online edition (email required)***

Name:

Organisation/Company:

Address:

Postal Code:

City:

Country

E-mail:

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/