Number 76, January 2009

# ERCIM NEWS

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

## Special theme:

# THE SENSOR WEB

## Also in this issue:

*ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. 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 per copy. Subscription is currently available free of charge.*

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



***Contributions***
*Contributions must be submitted to the local editor of your country*



***Advertising***
*For current advertising rates and conditions, see http://ercim-news.ercim.org/ or contact peter.kunz@ercim.org*

***ERCIM News online edition***
*The online edition is published at http://ercim-news.ercim.org/*

***Subscription***
*Subscribe to ERCIM News by: sending email to en-subscriptions@ercim.org or by filling out the form at the ERCIM News website: http://ercim-news.ercim.org/*

***Cover photo:***
*Synchronization of sensor data in team sports allows team-based queries with a potential to discover new knowledge from sensed data. Photo: AFP*

***Next issue:***
*April 2009, Special theme: Future Internet Technology*

*Barry Smyth*
*CLARITY: The Centre for Sensor Web Technologies, Ireland.*

# The Sensor Web: Bringing Information to Life

There is no doubting the impact that the information revolution has had on our everyday lives. How we learn, work, and play has been forever transformed by the always-on connectivity of the Internet. But this information revolution has largely been confined to the online world and, for many of us, we continue to interact with two very separate worlds: the physical world in which we live and the online world of the web. There has been a physical-digital divide. Every second of every day information is created through naturally occurring events in the physical world but these events go largely unnoticed and the information is lost. In the world of the Sensor Web this is set to change.

The catalyst for this change will come in the form of a new generation of cheap, reliable, and flexible sensor technologies, which will serve as new peripherals for the internet by bringing a whole new world of input data to the wider web. Accordingly, data will no longer just be generated from the keyboards and scanners of desktop PCs. Instead, these new sensor technologies will permit the sensing of diverse events in the physical world, from the traffic congestion in our streets to the pollution in our river systems, and from energy consumption in our cities to recycling in the home. Sensors that can be integrated with garments, and worn on the body, will permit the capture of physiological data as we exercise or recuperate. In short, this unique combination of sensors, software, and the Internet will enable new types of information services across a wide range of sectors from health and the environment to education, retail, and entertainment.

This is the vision of the Sensor Web. Its guiding principle is that better information helps people to make better decisions and that by harnessing the potential of the sensor web we can help people to live healthier, safer, and more productive lives. For example, decades worth of studies about energy usage in the home have shown that by simply informing people about the impact of their energy usage in real-time is sufficient motivation for them to cut their energy consumption by up to 15%.

Of course realising this vision is not without its challenges and as a result there is currently considerable research investment in Sensor Web technologies by funding agencies and industry players alike. Researchers at the new Centre for Sensor Web Technologies (CLARITY) are addressing these challenges head-on (see www.clarity-centre.com). CLARITY is funded by Science Foundation Ireland (SFI) and involves more than 60 researchers in University College Dublin, Dublin City University, and the Tyndall National Institute (TNI). The centre will focus on the development and application of sensor web technologies in areas such as personal health, environmental monitoring, and media and the research focus will include three different sensor modalities:

1. Sensing the Body – Focusing on the use of wearable technologies to capture physiological data form the wearer with key applications areas such as exercise, sports and personal health.
2. Sensing the Mind – Recognising the preferences of individuals and groups of users by mining sensed interaction data.
3. Sensing the Place – Focusing on the use of sensors in monitoring the world in which we live, with applications in environmental monitoring.

Within these modalities, the same challenges will arise: bridging the Physical-Digital divide, specification and deployment of adaptive middleware, and taming the information overload.

Addressing these challenges and developing practical Sensor Web solutions will, in the near time demonstrate the true potential of this technology to industry, government, and consumer alike. Indeed there are reasons to be optimistic that such technologies are accessible in the near-term. For example, modern mobile phone infrastructure provides an ideal communication network for the Sensor Web and mobile phones serve as ideal sensor platforms capable of capturing information from subscribers and their locale. Comprehensive location-sensing technologies such as GPS, for example, are now built in to most modern mobile devices and this is already inspiring a new wave of novel applications in which location sensing plays a key role. But all of this is just the beginning.

*Barry Smyth*

# EVOL 2008 - Fourth International ERCIM Workshop on Software Evolution and Evolvability

by Michel Wermelinger and Paul Wernick

***The fourth edition of this annual workshop took place in L'Aquila, Italy, 15-16 September 2008, under the auspices of the ERCIM Working Group on Software Evolution. The event gathered academics to present and discuss the state-of-the-art in research on software evolution.***

This year's workshop was the result of a merger between the ERCIM workshop on Software Evolution and the IEEE International Workshop on Software Evolvability. Historically, the ERCIM workshops have dealt with the practicalities of developing software that can be changed to meet evolving needs and the tools and methods by which this can be best achieved, while the Software Evolvability workshops have generally addressed issues at a more abstract level, advancing the understanding of the causes and effects of software evolution using means such as analogies and models from biology, the sociology of technology and other areas.

The workshop was supported by ERCIM co-located with the 23rd IEEE/ACM International Conference on Automated Software Engineering (ASE 2008). The organisers were Michel Wermelinger (The Open University, UK), Paul Wernick (University of Hertfordshire, UK) and Ciarán Bryce (INRIA Rennes, France).

The workshop attracted 20 participants from ten countries. The programme consisted of ten long and four short papers, one keynote talk, and a closing panel. To make the event a true workshop, the schedule had plenty of discussion time built in, session chairs were asked to prepare in advance some discussion topics, and participants were given the papers in advance. The papers covered a wide range of topics, from run-time changes to long-term maintenance, from formal models to empirical studies of the state of practice in industry. Moreover, joining the two workshop series proved to be successful, with some papers taking up the theme of 'bridging boundaries' and looking for new techniques and insights in other disciplines – in particular healthcare, natural language processing and epidemiology. The full list of papers and authors can be seen on the workshop's website (see link at the end). The proceedings will be published in the IEEE's digital library.

Jean-Marie Favre (University of Grenoble, France) provided in his keynote address under the title 'Past, Present and Future of Software Evolution: From Software-Now to Software-over-Centuries' a view of the size of the territory covered by software evolution studies, both in subject matter and time. He made an analogy between the evolution of informatics and the evolution of human societies and pointed to the importance of looking at software evolution at various time scales.

Finally, in the closing panel, Jean-Marie Favre, Massimiliano di Penta (University of Sannio, Italy) and Serge Demeyer (University of Antwerp, Belgium) presented and discussed with the audience their views on the challenges and opportunities facing the software evolution research community.

**Links:**
Workshop Web site: http://evol08.inria.fr
Working Group Web site: http://w3.umh.ac.be/evol/

**Please contact:**
For information on the workshop:
Michel Wermelinger
The Open University, UK
Tel: +44 1908 659 539
E-mail: m.a.wermelinger@open.ac.uk

For information on the EVOL Working Group:
Tom Mens
Université de Mons-Hainaut, Belgium
Tel: +32 65 37 3453
E-mail: tom.mens@umh.ac.be

# SERENE 2008 International Workshop on Software Engineering for Resilient Systems

by Nicolas Guelfi, Henry Muccini, Patrizio Pelliccione and Alexander Romanovsky

***The first workshop of the ERCIM Working Group on Software Engineering for Resilient Systems (SERENE) was held in Newcastle upon Tyne, UK, on 17-19 November 2008.***

The SERENE series of workshops, organized by the ERCIM working group SERENE, is an international forum for researchers and practitioners interested in advances in Software Engineering for Resilient Systems. SERENE 2008 was held in cooperation with ACM SIGSOFT, the International Conference on Computer Safety, Reliability and Security, and was attended by 42 participants. The technical programme contained two invited talks, three sessions of technical papers, two PhD forum sessions, a project session and an experience/industry session. The workshop organisers received 23 submissions (thirteen technical, three industry, five PhD and two projects), 17 of which were accepted for presentation (seven technical, three industry, five PhD and two projects). All papers have been peer-reviewed by at least three members of the program committee.

The first invited talk, entitled 'System Complexity, Dependability and Failures', was given by Brian Randell from Newcastle University, UK. Ralf Reussner of Universität Karlsruhe, Germany, delivered the second, on “Challenges and Results in Component Quality Certification”.

Organization of such a workshop represents a big challenge. We would like to acknowledge the help of the programme committee members, the additional referees, the organization committee members, and the support of the scientific, technical and administrative staff of Newcastle University, UK.

The workshop proceedings, ISBN 978-1-60558-275-7, are included in the ACM Digital Library.

SERENE 2008 was supported by ERCIM, FNR (Luxembourg National Research Fund), the ICT FP7 DEPLOY IP (on Industrial Deployment of System Engineering Methods Providing High Dependability and Productivity), LASSY (Laboratory for Advanced Software Systems, University of Luxembourg), School of Computing Science, Newcastle University, and the ICT FP6 ReSIST (Resilience for Survivability in IST) Network of Excellence.

**Links:**
http://serene.uni.lu/tiki-index.php?page=Serene2008Overview
http://serene.uni.lu/tiki-index.php

**Please contact:**
Nicolas Guelfi
ERCIM SERENE WG coordinator
Luxembourg University
E-mail: Nicolas.Guelfi@uni.lu

# Interlink Consolidation Workshop

by Dimitris Plexousakis

***The ERCIM-led Future and Emerging Technologies (FET) Coordination Action Interlink (International Cooperation Activities in Future and Emerging ICTs) held its consolidation workshop in Cannes, France on 12-14 November.***

Interlink is a two-and-a-half year endeavour aimed at identifying new challenges for basic research in three areas of strategic importance, namely Software-Intensive Systems and New Computing Paradigms, Ambient Computing and Communication Environments, and Intelligent and Cognitive Systems. The workshop brought together all three working groups with the goal of finalizing proposals on strategic research directions, related disciplines and fundamental problems that need to be addressed within the selected areas.

The first day of the workshop began with an overview of the aims and structure of the workshop by Dimitris Plexousakis, FORTH. This was followed by a presentation by Interlink's project officer, Walter van de Velde from the European Commission (EC), on the collaborative research opportunities arising within the programme of the FET unit of the EC in forthcoming years. The workshop also featured a number of invited talks. Joseph Sifakis (CNRS Grenoble, ACM Turing Award 2007 recipient) delivered the keynote address on the verification of complex software systems. The coordinators of the three thematic areas gave an overview of their respective research areas and group proposals, setting the context for the Thematic Group invited talks. Stefan Jaehnichen (Fraunhofer Gesellschaft) talked about challenges arising in grid computing and computing clouds. Ted Selker (MIT Media Lab) presented his research and ideas on context-aware computing and Henrik Christensen (Georgia Tech) elaborated on the study of artificial cognitive systems.

For the second day of the workshop, the three groups worked in parallel with the aim of identifying major challenges where significant added value is expected to be gained from worldwide cooperation. The groups also deliberated on the scientific, technological and societal impact of the research proposals as well as on knowledge-building mechanisms, research practices, innovation mechanisms and programmes for sustainable cooperation. On the third day, representatives of the Working Groups and the scientific steering committee discussed cross-thematic challenges arising at the intersections of the thematic areas. The group identified the vision of 'smart cities' (multiscale smart spaces) as a unifying research theme in which major challenges relating to the individual areas arise most prominently. The workshop concluded with presentations of the research proposals of the Working Groups as well as of the cross-thematic challenges.

*Interlink workshop participants.*

**Links:**
http://www.ercim.org/interlinkworkshops/
http://interlink.ics.forth.gr.

**Please contact:**
Dimitris Plexousakis
ICS-FORTH, Greece
E-mail: dp@ics.forth.gr

# New ERCIM Working Group on Data and Information Spaces

by Ingeborg Torvik Sølvberg

***18 September, 2009 was the kick-off date for the new ERCIM Working Group on Data and Information Spaces (DIS). The objective of DIS is to build a strong network of researchers in the fields of information repositories and digital libraries and their interoperation.***

DIS will focus on issues relating to digital libraries, and the storage, preservation and curation of all types of data, including primary and scientific data. Intelligent information management and interoperability and scalability will be addressed. The development of methods and tools to make digital, cultural and scientific content available, searchable and accessible to all kinds of users is of special interest.

Eleven ERCIM organizations have registered for the WG, suggesting the topic is relevant and timely for ERCIM members. At the kick-off meeting, eight organizations were represented by thirteen individuals. Among the issues discussed was the name of the WG. It was decided to change it from Working Group on Digital Libraries to Working Group on Data and Information Spaces, DIS. A draft of the work program for the two years was discussed. The WG intends to initiate research projects in areas that are of interest to its members, to arrange workshops and to encourage exchange of researchers and students. The next meeting is the open Workshop that will take place in Paris on 27 May, 2009 in conjunction with ERCIM's 20th anniversary meeting. Scientists interested in participating in the ERCIM DIS Working Group should contact the coordinator.

**Link:** http://wiki.ercim.org/wg/DIS/

**Please contact:**
Ingeborg Torvik Sølvberg
ERCIM DIS Working Group chair, NTNU, Norway
E-mail: ingeborg.solvberg@idi.ntnu.no

# ERCIM office in Brussels

ERCIM opened an office in Brussels in December 2008. As requested by ERCIM's Board of Directors, this will enable ERCIM to strengthen its links with European Commission key players. The new office is ideally located within the premises of our Italian member, CNR, in B-1050 Brussels, rue du Trône 98, right in the centre of the European quarter. ERCIM is now able to offer a hosting environment for ERCIM staff and ERCIM project meetings in Brussels. The office is headed by Pierre Guisset, formerly head of CETIC.

**Please contact:**
Pierre Guisset, ERCIM office Brussels
E-mail: pierre.guisset@ercim.org.

# ERCIM at ICT 2008

***ERCIM had a major presence at ICT 2008 in Lyon, France, on 25-27 November. This bi-annual event, organised by the European Commission Information Society and Media Directorate General, attracted over 4,500 delegates and presented Europe's priorities for information and communication technologies (ICT) research, development and funding.***

Through stands and networking sessions related to ERCIM itself or to European projects in which ERCIM is participating, the consortium had its highest exposure ever at "Europe's biggest research event for ICT".

ERCIM was present on its own booth and on the stands of the European projects VITALAS (Video & image Indexing and Retrieval in the Large Scale), EchoGRID (European and Chinese cooperation on Grid), GridCOMP (Grids Programming with Components: an advanced Component Platform for an effective invisible Grid), EuroIndia (Euro-India ICT Co-operation Initiative), and Digital World Forum on Accessible and Inclusive ICT. Networking sessions were held for InterLink (International Cooperation Activities in Future And Emerging ICTs), EchoGRID and EuroIndia.

## Interlink

As a coordination action supported by the EC IST Programme, InterLink held a networking Session entitled 'International Cooperation Activities in Future and Emerging ICTs'. The aims and results of InterLink were presented along with an outlook on opportunities for future international collaboration. The results presented were the new research proposals that came out of the InterLink consolidation workshop held in Cannes in November (see article on page 7).

## Euro-India

The Euro-India (Bringing European & Indian Perspectives for future ICT Co-operation) played a central role at ICT 2008 by holding a networking session entitled "Euro-Indian ICT Cooperation Gateway". Chaired by Hilary Hanahoe, Trust-IT Services, and opened by Alvis Ancans, European Commission, International Relations Unit, DG Information Society & Media who illustrated the EU-India ICT R&D cooperation, the session attracted more than 50 participants.

After a brief presentation of the India Mentor project (Mentoring Indian IT organisations in the participation in the ICT programme of FP7) by Mauro Bianchi (Teseo Sprl, India Mentor project coordinator) and the EuroIndia initiative by Mogens Kuehn Pedersen (Copenhagen Business School, EuroIndia Project Coordinator), the session continued with keynote speakers on different topics. Manesh Kulkarni, from CDAC Pune (India) presented the Digital Libraries projects in India. Prof. Dilip Kanhere from the centre for modelling and simulation, Pune University (India) explored distributed computing. A. Prabaharan from Jawaharlal Nehru University; New Delhi (India) gave an overview on ICT for learning. IPv6, a topic of huge interest was tackled by T R G Nair, director


*Euro-India booth.*


*Joint EchoGRID, GridCOMP and BRIDGE booth.*

of research and industry from the Dayanand Sagar Institution (India), while software and services in India were presented by Prashant Shukla, COO from CMC Ltd (a TATA enterprise). Finally, Krishna Lakshmi Narasimhan, vice president of Cranes Softwares (India) gave an insight into the successful Indian experience of the AEE (acquire-enhance-expand) model for intellectual property identification, development and commercialisation.

The session ended with intensive discussions between the participants interested in developing or increasing cooperation with India and interested in understanding more about EU-Indian ICT collaboration. The participants had the opportunity to continue their discussion and to learn more about the India Mentor and EuroIndia projects visiting the joint booth located in the 'international village'. Visitors viewed a rolling presentation with quotes on innovation by a select group of experts from India. The presentation highlighted India's leadership in specific ICT areas, and its potential to invent, innovate and make an impact through collaborative developments with Europe.

Delegates were also invited to participate in the series of events which will be held in India in January, such as 'Information Days' (Mumbai, 20 January 2009; New Delhi, 21 January 2009) and the First EuroIndia international conference in New Delhi, 22-23 January 2009. The conference will take place during the 'Indian R&D 2009: ICT Innovation' event, organized by the Federation of Indian Chambers of Commerce & Industry (FICCI). This multi-faceted event featured keynote talks, parallel sessions on ICT topics of interest to Europe and India, and an exhibition showcasing technological achievements in both India and Europe. EuroIndia conference will include a session presenting the main outcomes of EuroIndia's initial technology road-mapping and a guide on how to get involved in European Commission-funded projects as well as presentations in all sessions.

http://www.euroindia-ict.org/

## EchoGRID and GridCOMP

ICT 2008 offered a prime opportunity for EchoGRID and GridCOMP projects to build on their achievements and spearhead synergies with a number of EC-funded projects and technology platforms. The joint booth of EchoGRID, GridCOMP and BRIDGE (Cooperation between Europe and China to Develop Grid Application) on "EU & Chinese Co-operation on Grids, and ProActive/GCM Demonstration", at the 'International Village', demonstrated the benefits of Grids across a range of commercial and public sectors with live demos from Bridge and GridCOMP showcasing cutting-edge technologies developed with partners from China and other world regions, BRIDGE presented three application scenarios designed to demonstrate co-operative design, simulation and data access between European and Chinese partners, with special emphasis on the interoperability approach to combine heterogeneous Grid infrastructures, as well as the workflows in the Bridge application domains. The networking session on Thursday 27 November highlighted current and future developments and areas for future cooperation of interest to both Europe and China.

The GridCOMP ProActive/GCM (Grid Component Model) demonstration tackled the increasing need for both improved business performance and optimal energy efficiency now facing enterprise IT. GridCOMP has developed an open-source reference implementation in the ProActive parallel suite library that enables businesses to globalise their IT infrastructure while lowering costs and accelerating applications.

Business intelligence often involves data mining and ETL (Extract Transform Load): relevant information from enormous files are analysed, extracted, transformed and loaded into data warehouses in order to support decision making processes. GridCOMP has enhanced these data-mining and ETL processes and has delivered a faster response-time capability for different departments in a Telco-company scenario.

The demonstration showed how the GCM provides an easy way to use deployment framework providing interoperability with a large set of Grid schedulers and middleware. This is one of the industrial use cases developed within GridCOMP by GridSystems (http://www.gridsystems.com/).

Visitors also learnt about the achievements of EchoGRID in connecting actors from EU and China through a series of events in both regions to deliberate top-level challenges for new computing paradigms and define research priorities moving forward. In particular, the EU-China cooperation workshop on 28 October 2008 in Beijing, showcased successful partnerships between China and Europe in Grids, software and services, and e-infrastructures, with the aim of fostering future collaborative developments to address new

challenges. The workshop brought together around one hundred European and Chinese representatives, including the European Commission and Chinese Ministry of Science and Technology officials.

The networking session held on Thursday 27 November 2008, was entitled “Priorities for Future Research on Grids: EU & International Perspectives”. This Session looked at different research agendas around Grid from multiple yet complementary perspectives, with the future evolution of distributed computing and software and services firmly in mind. Each of the hosting projects, EchoGRID, NESSI-GRID (Network European Software and Services Initiative - GRID), CHALLENGERS and 3S (European Community for Software & Software Services), gave a brief talk on the main outcome of their road-mapping activities and white papers. BRIDGE then took the floor to present current achievements in Grid developments, paving the way for discussion.

The session brought into focus future research priorities for distributed computing and software and services defined by EU road-mapping projects both at a European and international level, featuring success stories on EU and Chinese co-operation on grids.

EchoGRID: http://echogrid.ercim.org/
GridCOMP: http://gridcomp.ercim.org

## Digital World Forum

The Digital World Forum EU project staffed a booth in the 'International Village' of ICT 2008 and issued a press release announcing a public workshop on the role of mobile technologies in fostering social and economic development in Africa in Maputo, Mozambique, on 1-2 April 2009. In this workshop, participants will explore ways to fulfill the potential of mobile phones as a platform for deploying development-oriented ICT services towards the poorest segments of populations in developing countries, with an emphasis on the African context. DWF also participated in the networking session on "Strategic cooperation opportunities with sub-Saharan Africa".

DWF is an FP7-funded project focusing on the use of ICT to leverage economic development in Africa and Latin America. The project explores ways in which advantage can be taken of the new model of low-cost technologies in broadband infrastructure and devices, as well as the explosion of mobile telephony to bridge the digital divide and help connect the hitherto unconnected.

http://www.digitalworldforum.eu/
http://www.w3.org/2008/11/mw4dafrica-pressrelease

## VITALAS

The VITALAS project (Video and image Indexing and reTrievAl in the LArge Scale) project demonstrated at ICT'08 the first release of its cross-media search system which combines semantic search, textual and visual concepts search, and visual similarity search on a corpus of 10,000 annotated images provided by the Belgian news agency Belga. The system was previously presented to a large audience during the the Networked and Electronic Media Summit in St Malo, France, 13-15 October 2008

The next release of the prototype, VITALAS V2, available in February 2009, will integrate text, video and audio modules. This version will enable audio queries and visual queries, analyse of non-textual content in such a manner that textual annotation from audio and visual content analysis can be automatically predicted and generated.

http://vitalas.ercim.org/

*VITALAS at ICT'08.*

## Launch of ERCIM Innovation

The first edition of 'ERCIM Innovation' was successfully launched during the ERCIM cocktail reception attended by some 200 people, including 30 representatives from the European Commission. This brand new magazine aims at communicating ideas and innovations from ERCIM members in order to facilitate their uptake by business and industry. In short, the purpose of this magazine is to foster innovation in Europe thanks to closer relationships between research, academia, standardisation, industry, and investment. This kind of activity demonstrates that ERCIM is not only the European network of research but also a unique European network of innovation. The content for ERCIM Innovation is provided by our new network of technology/knowledge transfer & innovation experts working in our member institutes, with additional articles of interest contributed by external innovation experts. The magazine is available for download from the ERCIM Web site. Printed copies are being distributed to 9,000 people throughout Europe and can also be requested from catherine.marchand@ercim.org.


*The first edition of 'ERCIM Innovation'.*

# Engineering Secure Complex Software Systems and Services

Executive Summary of the European Commission-ERCIM Seminar on ICT Security

***ERCIM and the European Commission jointly organised a Strategic Seminar on "Engineering Secure Complex Software Systems and Services". The seminar was held in Brussels on October 16th 2008 and is the result of an effort of ERCIM, its Security and Trust Management Working Group, and the European Commission's DG INFSO Unit F5 "Security".***

The seminar aimed at collecting the relevant academic and industrial expertise in secure software engineering (shortly, SSE) and at linking it with industry's best practices in the field. As the Information Society continues to develop, the security of its supporting ICT infrastructures will grow in importance. The need for assurance of software systems and services demands a set of novel engineering methodologies and tools in order to ensure secure system behaviour. There is clearly the urgency and, actually, the opportunity for exploiting synergies of advanced research approaches with industrial best practices in order to reduce the gap between theory and practice.

The specific objectives of the seminar were:
- to present the best practices applied in industry and to discuss latest progress on key R&D initiatives
- to encourage the dialogue and promote collaboration between scientists and industrial players
- to identify future key research challenges, in particular in the context of the evolution towards the Future Internet.

This report briefly describes the main findings of the seminar, which was attended by more than 60 stakeholders from industry and academia. The full seminar report, agenda, individual presentations and list of participants are available at http://www.ercim.org/activity/strategic_seminar.

## 1. Industrial Best Practices and Perspectives

The first panel of the seminar addressed industrial best practices in the field and future perspectives. The panellists also discussed about IT frameworks, models and tools required for improving the development of secure software throughout its lifecycle; creating a sound business case for security; promoting software assurance and measurability and testing procedures for auditing and security compliance purposes; dealing with the increasing complexity of IT systems; and, education, training and awareness initiatives.

*Best practices*

In ever changing and global markets, software companies are continuously developing and improving their procedures and tools for embedding security in their software systems and services. A rich set of best practices is now around in terms of documents and guidelines that ask for strict development process control, supervision, or review. Recently, joint corporate initiatives on secure software were also launched. They clearly demonstrate the great interest that major industrial players and private and public organisations have in cooperating in this field by sharing and promoting pragmatic approaches and proven software assurance practices. Automated support for best practice enforcement and the ability to reason about the business impact of security are key issues to manage security related efforts in an economically feasible way.

*Novel IT frameworks, models and tools during all phases of the software lifecycle*

Software security should be an integral part of every phase of the software lifecycle (ie, from design to deployment, monitoring and auditing). The existence of common IT development and execution frameworks enforces the use of best practices and fosters collaborative work towards their further improvement for achieving higher levels of secure software. Formalising and describing how the many possible processes and their security requirements have to be organised into an application or system is essential to software industry. Modelling tools could provide the right abstract schemes to make possible the description and assessment of alternative scenarios for achieving a balanced secure software solution. Furthermore, industry needs IT tools that support security in the software that it produces or uses and that are platform- and programming language-agnostic. In fact, industry requires tools that encapsulate specialised knowledge by translating underlying theoretical foundations into concrete secure software development practices. Such tools have to be well integrated into development environments and be easy to use by non-experts.

*Creating the business case for security*

Despite the accrued interest of industry on SSE practices, overall, IT security has to compete with several other investment priorities. With squeezing IT budgets and ever-shorter times to market, how much do managers need to spend on IT security to achieve enough security and when secure is secure enough? Understanding the value that investments on secure software can add through the product value chain is

## Strategic Seminars

With the seminar on "Engineering Secure Complex Software Systems and Services" the ERCIM Board of Directors embarked on the initiative to organise a series of annual strategic seminars on current topics within ICT and Mathematics, responding to the need for closer ties with the European ICT industry.

The rationale behind this strategic decision is manifold: to enhance the impact of research taking place within ERCIM institutes and Working Groups by actively disseminating results towards industrial stakeholders; to expose researchers to ongoing research activities with an industrial take-up potential; and to help bridge the gap between research and industrial practice.

http://www.ercim.org/activity/strategic_seminar/

vital for business and IT managers taking decisions on spending money on security. Specifically, managers need to understand how much risk their company is ready to take for a given threat and manage that risk accordingly.

*Dealing with assurance, measurability and testing*
Understanding the value of security and assessing and managing risks implies putting in place an appropriate set of "controls" at different levels, business, technology or processes. Such a control framework would allow prevention of vulnerabilities and monitoring compliance with internal or external security requirements, including legal compliance. That requires, however, putting in place an appropriate set of independent measurement and testing procedures for all phases of the software lifecycle as well as metrics for collecting data, auditing performance and, ultimately, proving/ensuring security by measuring it.

*Dealing with increasing levels of complexity of software systems*
Presently, the complexity is rapidly increasing when moving from the secure engineering of isolated application components to that of software systems that mix various infrastructure resources with application functionalities. Such software systems are usually built incrementally resulting in "systems of systems" with functionality often different from what their underlying components were designed for. Moreover, they increasingly rely on real-time dynamic composition involving third-party software components and services. Under these circumstances, achieving secure systems and secure software products is a huge challenge and key business success factor.

*Promoting education and awareness*
Security conscious and well-educated software architects and software developers are needed together with more investments on higher-level education, professional and on-the-job training. Dedicated awareness creation initiatives would also permit to stress the importance of secure software within managers, software architects, programmers and users.

## 2. Research Advances and Perspectives

The second panel of the seminar focused on promising research directions for engineering secure complex software systems. It addressed the following topics: security requirements engineering; model-based techniques and automated tools for the development of complex secure software systems; methods for secure coding and programming; the recent advances on methodologies and tools for the verification and validation of specifications and code; and finally, the role of risk in the creation of secure "systems of systems".

*Security requirements engineering*
Several security weaknesses originate in the incomplete or conflicting nature of security requirements of software code. Specific expertise, methods and tools should be devoted to this task. For example, a step-by-step refinement procedure (eg, model-based requirements design) and automated tools would help security requirements engineers to improve the process from requirements elicitation to analysis and to track them during the subsequent software development steps. Also, mechanisms able to pass from negative-form requirements to more operational ones (as for functional requirements) should be envisaged. As a whole, security requirements engineering is an area where progress is possible and potentially useful in order to answer common software industry needs.

*Models for Secure Software Engineering*
The software development process needs several models to deal with domain specific aspects and to identify the correct security solutions to adopt. These models often have to be combined and refined in a way that ensures that the overall security of the final product is kept. Appropriate techniques to pursue here are model-driven design, security patterns, and case modelling and analysis of "uses" and "abuses". Process description and model checking techniques could be used to validate specific solutions at a given design stage, eg for validating requirements. Design techniques should involve component-based approaches allowing modular verification – compositionality is in fact a major security challenge related to the scalability and inherent complexity of ICT systems. Another challenge to deal with, from a security point of view, is dynamic change of systems and code and dynamic evolution of system functionalities.

When applicable, formal methods seem to be able to guarantee an increased robustness of software. Today, the high cost of applying them is an impediment to their larger industrial deployment. Therefore, one of the research directions with major impact would be to embed formal methods in automated development tools in a transparent way for the user. Finally, methods for measuring the trustworthiness of the software systems, is yet another area of importance for industry where major research efforts are necessary.

*Language-based security*
Language-based security is regarded as the backbone of secure software engineering. Indeed, language-based security techniques and specific type systems allow verifying, at compile time, the absence of (certain) vulnerabilities and constrain the run-time execution of applications. In fact, they move the burden of ensuring the security of the final code from the application programmer to the programming environment developers. Further progress is expected from several ongoing efforts aimed at embedding information flow management techniques in programming languages such as Java, or at embedding security mechanisms in Business Process Execution Languages used for composing complex services. A promising research area is developing techniques for proving complex properties of cryptographic algorithms as well as provably correct implementations.

*Advances in security verification and validation*
Several rigorous techniques have been developed for checking system specifications, such as model checking and theorem proving. However, there are still several limitations that must be addressed for their wider deployment in industry. Relevant research issues include addressing their scalability and coping with the ever-increasing complexity of software-intensive systems. In addition, one needs to take into account the uncertainty about the behaviour of the system components (eg malicious software) as well as external threats. Overall, more research efforts are necessary to make security

verification and validation tools usable in practice during software development at industrial scale.

*Advances in risk assessment for systems of systems*
Risk is a crucial notion in security and its role in the design of complex systems of systems needs to be further investigated. Issues to address here include assessing the complexity and the (cyclic) interdependencies inherent in ICT systems, often composed of several parts developed by different parties; and, assessing risks linked to changes in the lifecycle of systems of systems through, for instance, compositional risk-assessment methodologies. Embedding risk in an explicit manner in all the steps of the software development lifecycle could help to reduce the cost and make the improvements in software engineering more concrete.

## 3. The Way Forward

The last panel considered the findings from the two first panels and brought up some additional aspects related to: (a) enabling methodologies and tools for building secure complex systems and services; (b) software liability aspects; and (c) standardisation, education and other relevant issues for the field.

*Enabling methodologies and tools for building secure complex software systems*
Security engineering and software engineering methodologies and platforms should be integrated. The general (wrong) perception is that software engineering is dealing with construction of correct software, while security engineering is dealing with the deployment of software. The software architecture should be the starting basis. Security, manageability and scalability should be the main drivers for the software architects. Industry also needs usable and efficient methodologies and tools that automate the security of software code. Formal methods have proven to be useful for checking security specifications but not really software implementations. It is therefore urgent to undertake further work for bridging the gap between fundamental theories and pragmatic approaches for industry to use.

Industrial software is often built on top of legacy systems and/or is outsourced. This calls for tools for verifying the security properties and performance of legacy systems and/or third party software. The composition environment should permit to control the security properties of composed software both at the design phase and dynamically, at run time. Compositionality is a big challenge. Even if a software system is built from individually trusted components, the overall system may not be trusted. Modular verification of smaller modules may prove to be a good solution in large complex systems.

*Software liability*
For the moment, software companies in general and those companies in particular offering packaged software services or Service Oriented Architecture (SOA)-based applications and services are not liable for the likely damages they may cause due to software vulnerabilities of their products. As liability may change with time, it is important for companies to adopt best practices quickly. Should software companies become liable, they would need to become in full control of all the products, applications and services they sell, including all the underlying technology components supporting such products. A prerequisite for solving software liability is solving the compositionality problem.

*Standardisation, education and other relevant issues*
Currently there is a lack of sufficient standardisation in software security. In some cases, clear specifications are available at a certain level of abstraction, but implementations of standards are often not completely in line with these specifications. Robust tools for testing and validating such implementations are necessary.

Often there is a gap between the methodologies that secure software engineers are taught in Universities and the knowledge they need when working in industry. A closer and more productive cooperation is required between industry and academia in order to produce curricula dealing with both foundational knowledge principles and industrial reality.

## 4. Concluding Remarks

The significant participation of both industry and academia representatives at the event is showing the relevance of the topics addressed. Industry is showing sufficient motivation for adopting best practices in the SSE field and the scientific community can already bring several methodologies and tools. Targeting specific priorities as some of the ones identified in this report would certainly help to close the gap between foundational and practical work. Security and software engineering need also to be integrated in one coherent framework. As the complexity of ICT systems increases, easy-to-use software tools that encapsulate highly intensive specialised knowledge need to be developed through research and industrial partnerships. In order to ease this process, industry and academia should share similar expertise and adopt the same language and terminology.

Raising current levels of education and awareness in the field is another main issue emerging from the discussions held. Finally, special attention must be given to new forms of IT infrastructures such as cloud computing, “the internet of things” or, more broadly, the Future Internet, that bring new challenges for secure software as well as new opportunities for industry and business organizations.

**Links:**
*ERCIM Strategic Seminars:*
http://www.ercim.org/activity/strategic_seminar/

*ERCIM WG on Security and Trust Management:*
http://www.iit.cnr.it/STM-WG/

*European Commission's DG INFSO Unit F5 "Security":*
http://cordis.europa.eu/fp7/ict/security/home_en.html

**Please contact:**
Dimitris Plexousakis, FORTH-ICS, Greece
E-mail: dp@ics.forth.gr

Fabio Martinelli, IIT-CNR, Italy
E-mail: Fabio.Martinelli@iit.cnr.it

Thomas Skordas, European Commission
E-mail: Thomas.Skordas@ec.europa.eu

Introduction to the Special Theme

# The Sensor Web: Bridging the Physical-Digital Divide

by Mark Roantree and Mikko Sallinen

One of the truly multidisciplinary research efforts involving computer scientists revolves around the topic of sensor networks. It brings together chemists who develop the sensors, engineers focusing on wireless platforms and other hardware components, and the computer scientists who develop the services, knowledge layers and middleware. In many cases, research must also include the knowledge workers associated with the specific domain, many of whom are represented in the articles in this issue of ERCIM News. In almost all cases, some aspect of the research will seek to create a bridge or bidirectional channel between the physical world of the planet, its people and the sensors, and the digital world of computers and their software applications.

The emergence of the Sensor Web concept is due to the proliferation of physical devices that are accessible through the internet and thus, act as an extension to the World Wide Web. Through new hardware peripherals, connected directly to the Web, automatically interpreted, integrated and transformed for human interaction, querying and mining, we create the Sensor Web.

The Sensor Web provides a platform for new ideas and applications for different domains. However, each application domain has its own unique characteristics and the concept of a general platform can be developed only for laboratory tests. As a result, development tends to focus on proprietary solutions to meet a varied set of requirements.

The breadth of research in the Sensor Web domain is demonstrated in the articles in this issue of ERCIM News. Wireless networks are necessary to connect to sensor devices that may be physically unreachable. While sensor data is often analysed after sensing has stopped, wireless networks are essential if we are to perform live queries of sensor output, and adapt the behaviour of the sensor in real time. Toolkits for maintaining sensor networks, together with standards for processing and managing sensor data, assist in building more powerful and robust networks. In addition, new technology for developing and integrating smaller nodes enables measuring devices to be placed in a far wider range of products. In an increasing number of applications and projects, data generated by sensor devices is of a confidential nature, perhaps in areas such as personal health or body networks. This requires the appropriate research effort into security for the data transmitted by the many sensor devices inside the networks. Ubiquitous systems will also provide significant data volumes and challenges for the Sensor Web. Similarly, environmental monitoring needs to be continuous, integrated and without loss of data, requiring the specification and deployment of software services for the Sensor Web. Personal health (or pHealth) networks are emerging in many research projects and industrial applications. Wearable sensors transmit a variety of sensed readings from human participants, which are harvested and undergo semantic interpretation to allow domain specialists to make informed decisions on the health and increasingly the performance of individuals in sporting environments. There are several commercial products in this field that enable researchers and companies to develop more advanced solutions for the market. Applications of the Sensor Web covered in this issue include exploration in oil and gas fields, multimedia sensing, life-logging of human actions and interaction, and environmental hazards, demonstrating the multi-disciplinarity of Sensor Web research and highlighting the need to bring expertise from different backgrounds together.

So what issues arise from the articles presented here? There is evidence of a large number of sensor networks in different disciplines, as already discussed. They incorporate both small cheap devices and larger customized, proprietary and highly expensive devices. In general, the sole difference between them is accuracy. As the smaller, less accurate sensors become cheaper, they will quickly represent the significant majority of this device type on the Sensor Web. In many cases, more than one sensor will be required to support decision-making processes. This will demand synchronization and normalization of sensor feeds before integration takes place. While this presents problems, as highlighted in a number of the articles, the power of the Sensor Web is that it provides an infrastructure for harvesting the data. Historically, significant volumes of data generated by sensing devices have been lost, mainly due to a lack of computer scientists in the research project. This illustrates the gap between the physical and digital worlds.

The first step in developing a Sensor Web system is the construction of a simulation for the planned sensor network. This process becomes easier with time as the domains and environments are better understood. However, when designing a new architecture or software service for one of the layers in an architecture, or perhaps for a new domain, precise simulation is of considerable help prior to implementation into real components. A necessary requirement is that all layers are accurately modelled, otherwise the simulation will give misleading results. When this step is completed, the physical process of sensor deployment and sensor network construction can begin.

# New ERCIM Working Group on the Sensor Web

A new ERCIM Working Group on the "Sensor Web" was recently established. Representatives from seven ERCIM members (IUA, ICS-FORTH, ISTI-CNR, CRCIM, VTT, SARIT and NTNU), participated in the kick-off meeting on 19 May 2008, or expressed their interest in joining the Working Group.

## Objectives

The objectives of the Sensor Web Working Group is to promote and facilitate interactions between various R&D groups inside and outside ERCIM, in multidisciplinary themes relevant to the Sensor Web. The Working Group members cover a wide range of ICT skills (software engineers, information management and databases including information retrieval, wireless applications, networks, security and e-mobility, ambient and ubiquitous computing), and through their collaborators have a broad multidisciplinary base. Areas of interest to the group include both applied and basic research. Examples of deployment areas include personal health, environmental analysis, ambient intelligence, locomotive and large vehicle monitoring, military applications, deployment of personnel in toxic environments, and traffic analysis.

## Future Plans

The group intends to establish an ERCIM International Conference on the Sensor Web. This should be of a high quality, with the intention of raising the group's profile, providing a forum for discussing the meaning and scope of the Sensor Web, and attracting new people for future collaborations. The launch of a journal on Sensor Web has also been considered for the longer term. The Working Group is preparing for the next round of FP7 calls and intends to participate fully with ERCIM's Fellowship Programme.

ERCIM Working Groups are open to any researcher in the specific scientific field. Scientists interested in participating in the ERCIM Sensor Web Working Group should contact the coordinator.

**Link:**
http://wiki.ercim.org/wg/SensorWeb/

**Please contact:**
Gregory M. P. O'Hare
Sensor Web Working Group chair
The CLARITY Centre for Sensor Web Technologies, University College Dublin, Ireland
E-mail: gregory.ohare@ucd.ie

The first task for computer scientists in Sensor Web research is to ensure data crosses safely from the physical to the digital world, where it can be processed and manipulated to better inform us at to how to proceed in the many environments in which sensors now exist. In the past, environmental and personal health sensor networks have generated large volumes of data that were not captured in a digital format. For example, sports scientists have for some time been running sensor-based tests on athletes, in many of which data is recorded manually and on paper. Even complex sporting equipment such as speed gates that record the velocity of players as they move through various sections of a training course, will have their data transmitted wirelessly to a handheld device, which then has no means of transferring the data to a persistent storage mechanism. Eventually, sensor hardware will always provide a means of recording and transferring data but in the meantime, it is the role of data management researchers to devise a means for ensuring that this data is recorded electronically and stored in persistent, query-capable systems.

Analysis of the data and uncovering the essential issues from huge volumes of information is the next step. If data remains in the raw format generated by sensors, many knowledge workers will be unable to express the complex queries that are required to extract knowledge or make the assessments required to adapt the behaviour of sensors within the network. The next challenge for computer scientists is to convert the raw data into a usable format, preferably one that can be queried and updated by standard query languages such as SQL and XPath or XQuery. It is likely that the XML query languages will find widespread use, as data converted to XML has highly interoperable properties. This is crucial when integrating sensor data, both within a single sensor network and with the data generated by other sensor networks.

The final challenge is the identification of the most important issues in each application, eg closing the loop, controlling parameters, devices or actuators or giving instructions. These tasks should all be completed by accurate miniature sensors and nodes, wideband data communication, and the utilization of real-time control with minimum or zero power consumption. This final step demonstrates the path from the original sensor device through the engineering layers required for transmission of data, through the software services and human interaction, and finally back to the sensor, where the knowledge generated is used to make the sensing device more powerful and more accurate. The aspiration for the Sensor Web is that it should continue to evolve and address its limitations, so that the outcomes of the sensor age lead to an improvement in the planet's environment and the health of its citizens.

Please contact:
Mark Roantree, Dublin City University, Ireland
E-mail: Mark.Roantree@computing.dcu.ie

Mikko Sallinen, VTT, Finland
E-mail: Mikko.Sallinen@vtt.fi

# Sensor Networks in the Real World

by Steven D. Glaser and Tommi Parkkila

***At the Center for Information Technology Research in the Interest of Society, University of California, Berkeley, USA, we have been developing and deploying wireless sensor systems for ten years. We have focused on solutions to societal needs. As we ask more and more from our 'motes' and their low-power networks, we foresee important applications for sensor and control networks that will require a more powerful and flexible solution.***

The system we propose, AdapSys, is based on a fundamental unit that can perform at a very high level of abstraction - a multi-level controller and sensor hub that is completely software reconfigurable, including basic and ancillary functionality. In this scheme, each unit can act as a single complex controller as part of a locally controlled mesh, which in turn can be part of a wider distributed or hierarchical control network. All elements of this system consist of the same hardware, but have fundamentally fluid behaviors based on software adaptivity and reconfigurability.

We want to know what our structures are doing: structures in the big sense, from our bodies up to large industrial processes, airframes and buildings. This has traditionally been a troublesome and expensive problem. Recent improvements in sensors based on Micro-Electro-Mechanical Systems (MEMS) and in wireless technology have allowed the proliferation of wireless sensor networks, and these have completely changed what we can measure. Each element of the network is commonly called a 'mote' or smart sensor. Motes are combined into large networks that allow dense and detailed sensing. These networks move beyond the idea of a sensor as a single instrument measuring one thing, to a comprehensive system consisting of many small nodes working cooperatively. Engineering and science, however, remain captive to the traditional hierarchical embedded system. This experience has led us to devise a new monitoring and control appliance, each interacting in an organic network.

Here is an example of the current state of practice. During 2006 a mote network was designed, implemented, deployed and tested on the Golden Gate Bridge in San Francisco, in order to monitor its structural condition. Sixty-four motes were distributed over the main span and southern tower (see Figure 1), comprising the largest wireless vibration sensor network ever installed for structural health monitoring purposes. The spatially dense array resulted in an increase in effective signal-to-noise ratio compared to single, isolated, sensors, and most importantly allowed the higher modes, both vertical and torsional, to be analyzed easily and accurately.

## Deep Underground Science and Engineering Laboratory (DUSEL)

DUSEL is a large physics and engineering laboratory being constructed in the old Homestake gold mine in Lead, SD, USA. We are developing a deep in situ seismic observatory that will move us closer to the realization of rapid imaging of dynamical geo-processes at depth. More than 12 000 small-diameter (~ 65 mm) exploration holes exist throughout the mine, which we intend

*Figure 1: Mote antenna for the accelerometer package at the top of the South tower of the Golden Gate bridge.*

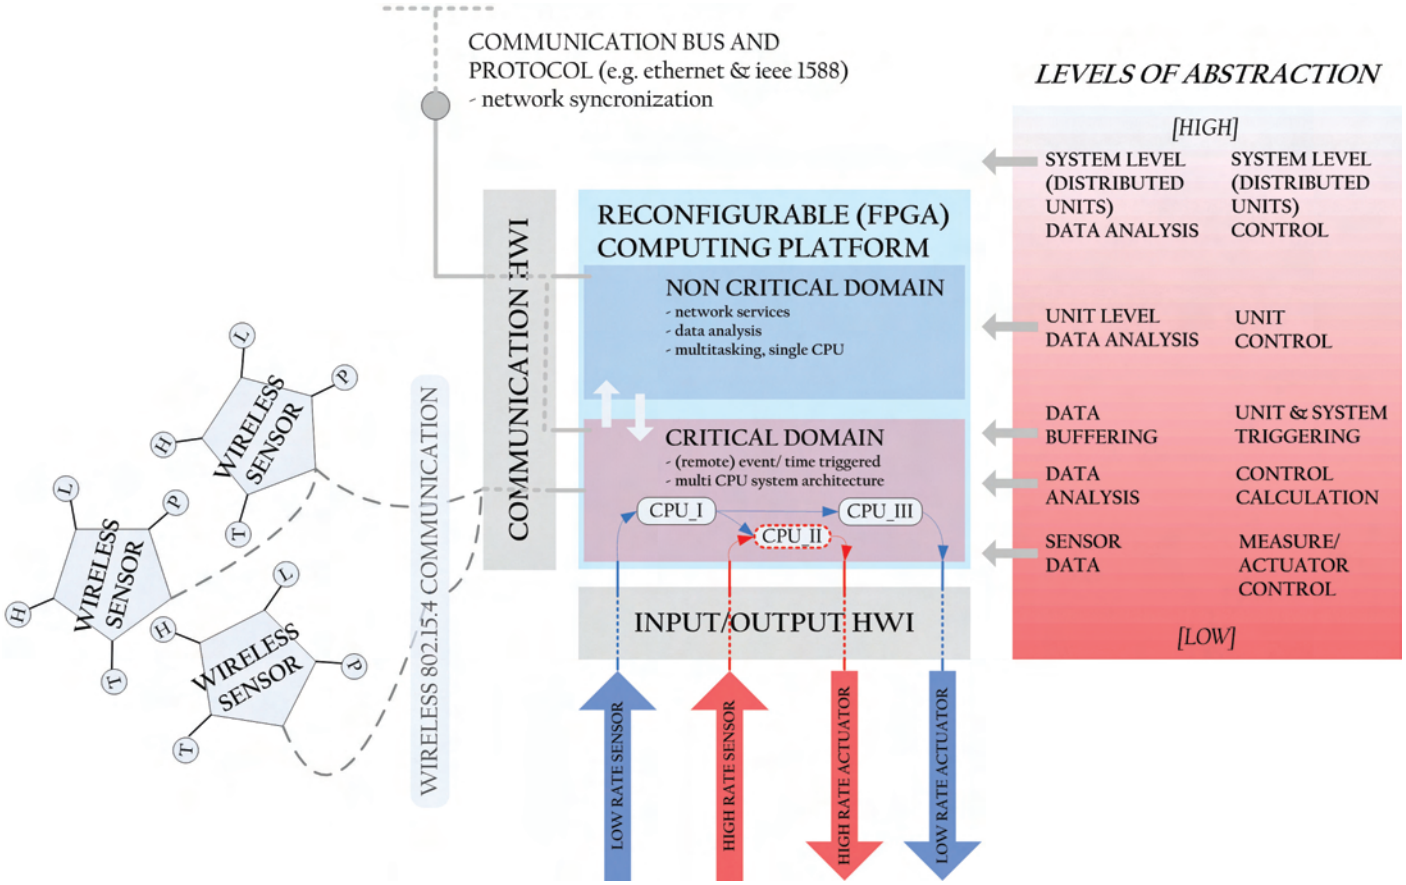


*Figure 2: The fundamental system unit of the AdapSys system., which includes options to act as a wireless sensor network.*

to use as multi-point monitoring probes. Any motion in the rock mass is thus surrounded by multiple receivers, which greatly constrains the inversion back to source movements. This solution led us to propose the AdapSys appliance.

## AdapSys

AdapSys is an elegant, straightforward, flexible and reconfigurable system comprised of Field Programmable Gate Array (FPGA)-based units. Each AdapSys unit is: (i) a real-time multi-channel data acquisition platform; (ii) a multi-sensor data aggregator; (iii) a remotely reprogrammable multilevel controller; (iv) highly portable; (v) distributed; and (vi) an embedded sensing and control network solution. We envision a group of AdapSys units controlling, say, a large paper mill, public conveyance systems, public safety equipment during a natural or man-made disaster, or even an array of wind generators with built-in nondestructive evaluation systems.

AdapSys uses a single FPGA to orchestrate and carry out the application demand through sensing, control and computation. AdapSys is a compound of flexible, reconfigurable, FPGA-based fundamental system units (see Figure 2). One of these units can execute the functionality of several microcontrollers through its multi-processor capabilities. New functionalities can be added to the system as parallel self-contained processor units inside the single FPGA chip. This allows the system to be incrementally upgraded in the field while allowing support of modular verification and certification.

Analyses of the DUSEL results along with past experience show that improved seismic arrays can be constructed from a string of accelerometer pods installed along a bore hole. We are currently prototyping down-hole sondes based on an Altera FPGA. The device has 24 input channels, both analog and digital, with a virtual real-time machine for each. Within the FPGA there are also real-time machines for the real-time clock, bus handling, and numerous control loops. All memory functions are handled seamlessly within the FPGA. Because the heart of the sonde is now software on the FPGA, there is little need for upgrades to entail physical hardware replacement; a completely new set of machines can be implemented by installing new software over the Web.

The AdapSys prototype is currently being assembled at the VTT Laboratory in Oulu. We are planning a joint research project with companies from the machine and automation industry in Finland, in order to test and refine the system in actual field situations.

**Please contact:**
Steven D. Glaser,
Center for Information Technology Research in the Interest of Society,
University of California, Berkeley, USA
E-mail: glaser@berkeley.edu
http://www.ce.berkeley.edu/~glaser

Tommi Parkkila
VTT - Technical Research Centre of Finland
E-mail: Tommi.Parkkila@vtt.fi

# MARWIS: A Management Platform for Heterogeneous Wireless Sensor Networks

by Gerald Wagenknecht, Markus Anwander and Torsten Braun

***A heterogeneous wireless sensor network (WSN) contains different types of sensor nodes. To operate such a WSN, we present MARWIS (Management ARchitecture for WIreless Sensor Networks). It uses a wireless mesh network as a backbone and offers mechanisms for visualization, monitoring, reconfiguration and updating program code.***

A wireless sensor network (WSN) may run different applications for different tasks, such as event detection, localization, tracking, or monitoring. Different types of sensor node are therefore required, and to handle heterogeneous WSNs with a large number of these different sensor nodes, a comprehensive management architecture is also necessary. We present MARWIS, a Management Architecture for heterogeneous Wireless Sensor Networks, which supports common management tasks such as visualization, monitoring, (re)configuration, updating and reprogramming. It takes into account the specific characteristics of WSNs and the restricted physical resources of the sensor nodes. These include battery life, computing power, memory, network bandwidth and link quality.

One of the main features of MARWIS is its hierarchical architecture. We divide a large heterogeneous WSN into smaller sub-networks, each of which contains sensor nodes of one specific type. A wireless mesh network (WMN) operates as the backbone and builds the communication gateway between these sensor sub-networks, the WSN and the Internet. Wireless mesh nodes perform the management tasks, and are controlled by a management station located in the Internet. A possible scenario is shown in Figure 1.

The use of a hierarchical architecture has various advantages. Sensor nodes, which are normally unable to communicate with each other due to incompatible radio chips, can be interconnected using wireless mesh nodes. Furthermore, dividing a huge WSN into smaller sensor sub-networks decreases the number of hops required to reach each sensor node. Specifically, each sensor node reaches the next wireless mesh node (which is the communication gateway) within three to four hops. This results in better communication performance with a lower round-trip time, lower jitter and less packet loss. A further advantage of using a WMN is that a new sensor node platform can be easily integrated into the heterogeneous WSN.

The architecture used to manage heterogeneous WSNs efficiently contains the following structural elements: one or more management stations, several mesh nodes as management nodes, sensor node gateways plugged into a wireless mesh node, and the heterogeneous sensor nodes. The management functionality is placed on the wireless mesh nodes, meaning the resource-limited sensor nodes have fewer management functions to perform, which in turn reduces memory and computation requirements. A user can perform management tasks using a management station, and this can be remotely located on the Internet.

Using a graphical user interface, the topology of the heterogeneous WSN with all the sensor sub-networks is visualized. The status information about every sensor node is monitored and displayed. This includes hardware features

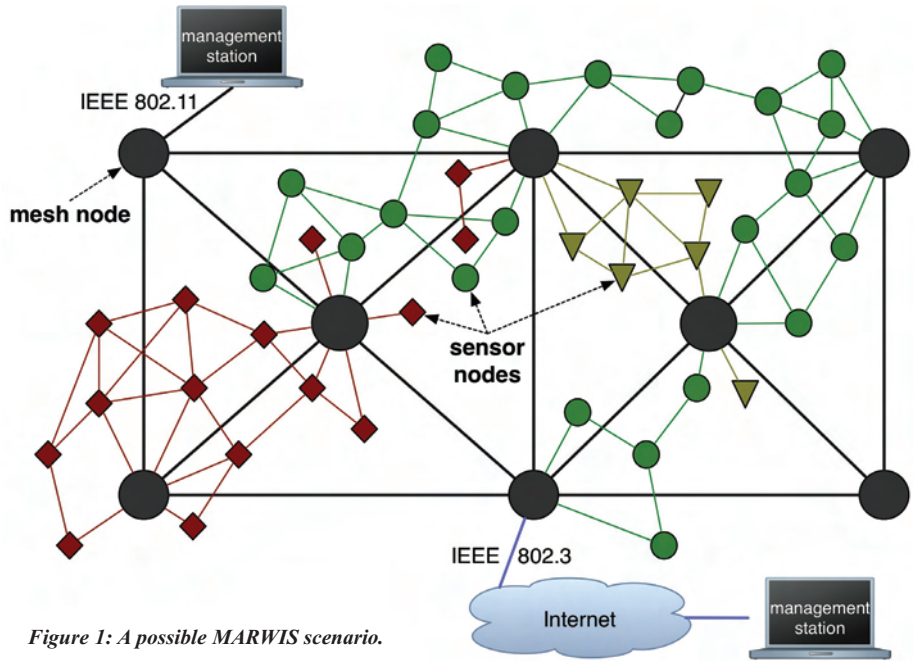


*Figure 1: A possible MARWIS scenario.*

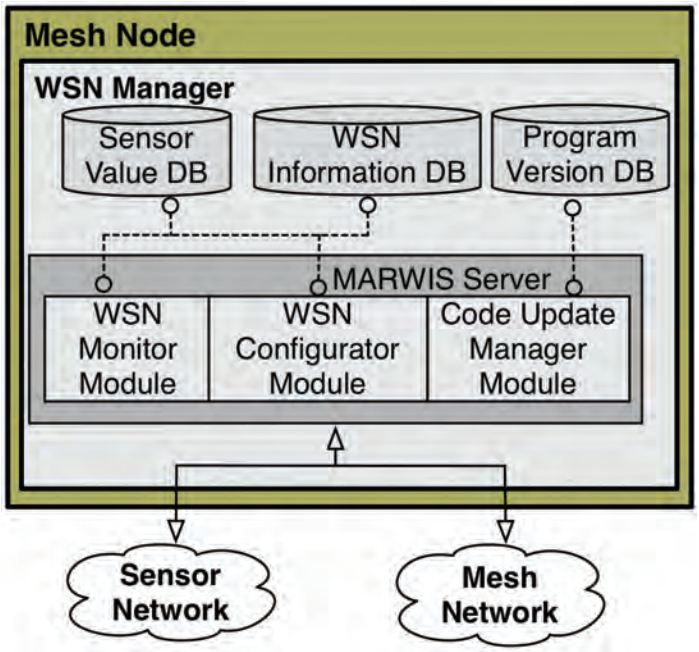


*Figure 2: The WSN manager provides the management functionality for the different sensor sub-networks. It consists of three databases and the MARWIS server with three modules.*

(micro-controller, memory, transceiver), software details (operating system versions, protocols, applications), dynamic properties (battery, free memory) and, if available, geographical position information. The applications running on the sensor nodes or network properties can be reconfigured using the user interface. Furthermore, updating and reprogramming the sensor nodes is a very important issue. In large WSNs manual execution of this task is unfeasible, and a mechanism to handle it automatically and dynamically over the network is required. Both the operating system and applications must be updated, either fully or partially.

The WSN manager located on the mesh nodes provides the management functionality for the different sensor sub-networks. It consists of three databases and the MARWIS server with three modules, as shown in Figure 2.

The WSN information database stores all information about the sensor nodes and the WSN, such as the topology (neighbours, addresses) and states of the sensor nodes (battery, memory). The program version database stores all versions of all programs for all platforms, which can be installed in the sensor nodes. Finally, the sensor value database stores all data measured by the sensors. To get information about the sensor nodes, first the databases on the relevant mesh node are queried. This means a direct connection to the sensor node is unnecessary, which saves time and energy. If newer information is required, the sensor node can be queried directly.

The MARWIS server contains three modules for the management tasks. The WSN monitor module connects to the WSN information database and to the sensor value database in order to handle requests from the management station. It also stores data coming from the sensor nodes into the databases. The WSN configurator module is responsible for the configuration tasks. It queries properties from the sensor nodes and stores them in the WSN information database. The code update manager module stores newly received program images (and related information) in the program version database and notifies the management station about available programs.

The Sensor Node agent is the complement of the MARWIS server and performs the management tasks on the sensor nodes after message exchange with the MARWIS server.

The architecture is currently being implemented and tested in a small real-world testbed. A small Linux distribution (kernel 2.6.14.6) is running on the mesh nodes; the MARWIS server is being implemented in C using sockets; the databases are managed with MySQL; the API for accessing the databases is implemented in C; and Contiki is running on the sensor nodes as the operating system.

**Link:**
http://www.iam.unibe.ch/~rvs/research/mancom.html

**Please contact:**
Gerald Wagenknecht
University of Bern, Switzerland
Tel: +41 31 511 26 36
E-mail: wagen@iam.unibe.ch

Markus Anwander
University of Bern, Switzerland
Tel: +41 31 511 26 34
E-mail: anwander@iam.unibe.ch

Torsten Braun
University of Bern, Switzerland
Tel: +41 31 511 26 31
E-mail: braun@iam.unibe.ch

# The IBM Mote Runner

by Thorsten Kramp, Michael Baentsch, Thomas Eirich, Marcus Oestreicher, Ivan Romanov and Alexandru Caraças

***Wireless sensor networks may well be the next big thing. Nevertheless, a fully business-process-integrated infrastructure for deploying large numbers of sensors and actuators requires a well-designed ecosystem. This should combine inexpensive devices with simple, bulletproof device programmability for easy integration and use by application domain specialists. The IBM Mote Runner system addresses this challenge with a high-performance, low-footprint middleware platform comprising a hardware-agnostic and language-independent virtual machine together with development and integration tooling to easily create and manage applications for open sensor and actuator networks.***

One-way, dedicated data-gathering IT networks such as those underlying, for example, a delivery tracking system, have shown the commercial value of real-time control of real-world components. Building on this, more generalized applications for wireless sensor networks (WSNs) are becoming increasingly apparent and significant in size and real-world relevance. Conceptually, the broadest application categories for WSNs involve environmental information, and provide a flexible communication and intelligence-gathering infrastructure that serves, for example, next-generation business applications by allowing them to directly tap into the ever greater number of digitally-enabled sensors and actuators that provide input to and control of their operation.

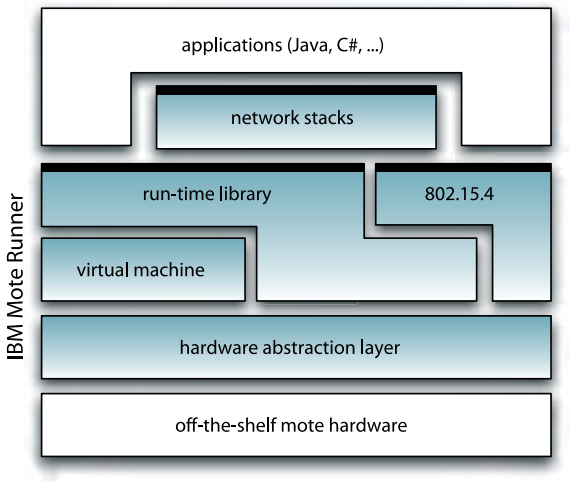


*Figure 1: IBM Moterunner architecture.*

To unlock this potential, however, two first-order problems must be addressed. One is cost: WSNs consist of many small computing elements that must be cost optimized. In this realm, cost takes the form of up-front investments in hardware and software plus any subsequent investments (eg for maintenance). Minimizing up-front investment requires minimum hardware cost, and this necessitates very efficient software running on the least expensive and most cost-effective off-the-shelf chips (or 'motes'). Minimizing subsequent investments translates into design requirements for minimum hardware interaction after mote deployment (eg for manual battery change or system reconfiguration).

The second problem is technological: a WSN runtime environment must not only be able to cope with the broad range of technical challenges imposed on WSNs but it must equally be accessible beyond the low-level functionality of individual WSN nodes. Here, 'accessible' refers to three things. First, it must be possible to dynamically configure and reconfigure the WSN in the field to deal with situations such as interrupted communication or WSN node failures. Second, it is necessary to secure the WSN in order that it may be considered a trusted source of information and reliable performer of actions in response. Third, the WSN must be well integrated into the larger infrastructure with which it cooperates. It must be generally programmable by domain specialists to solve domain-specific problems without deep knowledge of WSN technology and components. Only then are real-world solutions possible which link – while being easy to program and deploy – the physical world of sensors and actuators with business processes and applications. The result is the desired improvement in the responsiveness of transactions, enabling end-to-end process security and reducing cost by effectively using WSNs for data collection, pre-processing and autonomic feedback.

The IBM Mote Runner run-time environment for wireless sensor networks, currently under development at the IBM Zurich Research Laboratory, tackles these challenges in a holistic manner. At its core, Mote Runner provides a high-performance, resource-efficient virtual machine that is compatible with high-level languages and which shields portable applications from hardware specifics. It is designed to run on very small standard embedded controllers including low-power 8-bit processors, thereby reducing initial investment costs. Furthermore, it allows programmers to use object-oriented programming languages and development environments such as C# and Java to develop portable WSN applications that may be dynamically distributed, loaded, updated, and deleted even after the WSN hardware has been deployed, thereby reducing post-deployment and maintenance costs. All operations and communications can be cryptographically protected to establish a trusted execution environment. Figure 1 illustrates how this all fits together. Finally, Mote Runner WSN applications provide seamless integration with state-of-the-art back-end infrastructures by means of an event-driven process engine, which effectively bridges the gap to large-scale business and scientific applications without requiring deep technology skills.

**Link:**
http://www.zurich.ibm.com/moterunner

**Please contact:**
Thorsten Kramp
IBM Zürich Research Laboratory, Switzerland
E-mail: thk@zurich.ibm.com

# A Testbed for Sensor Service Networks

by Thomas Usländer and Kym Watson

***Working towards 'plug and measure' in sensor networks for environmental monitoring with Open Geospatial Consortium (OGC) standards, the SANY (Sensors Anywhere) project specifies an architecture for all kinds of fixed and moving sensors. This will allow both seamless plug-and-measure capability for sensors in the field, and sharing of information between sensor networks.***

The SANY project focuses on interoperability of in-situ sensors and sensor networks. This is done using both the standards and the on-going work of the OGC (in particular the Sensor Web Enablement suite of standards), OASIS (Organization for the Advancement of Structured Information Standards) and W3C. The SANY sensor service architecture provides a quick and cost-efficient way to reuse data from sensor and data sources that are currently incompatible. Data sources can include live sensor data, databases of archived data and model-based calculations.

The sensor service architecture and the service specifications have been made publicly available on the SANY project server, while the SANY specifications and best practice experience have been contributed to the OGC standardization work. The results are being tested in three innovative risk management applications covering the areas of air quality, marine risks and geo-hazards.

The Fraunhofer Institute for Information and Data Processing (IITB) has realized a testbed for sensors and services in order to trial the architecture and specifications. At the sensor network level, the ad hoc wireless ZigBee network is complemented by simulated sensor nodes, which measure properties such as temperature, humidity, illuminance and acceleration. The testbed is designed for experiments in a wide range of scenarios and scales, such as mobile sensors traversing several networks. The simulation is implemented as an application in LabVIEW (National Instruments), which has the additional task of configuring the ZigBee nodes. New sensor nodes (either real or simulated) are recognized automatically and registered in one of three OGC Sensor Observation Servers (SOS). The sensor values are then inserted into an SOS as they arise by measurement or simulation. The available network resources (observed features, sensors, services) are registered in a catalogue server along with metadata to support resource discovery by client applications. Clients can find, for example, information sources for a given region and observable phenomenon of interest.

Fraunhofer has produced in the testbed a special SOS known as a Fusion SOS, which is able to aggregate or fuse sensor data from several SOSs. The Fusion SOS queries the catalogue for available SOSs of the required type and then conducts a selected procedure to produce a spatial or spatio-temporal interpolation. The interpolation result is a so-called coverage, a function defined on a space-time grid of sampling points. The procedure takes the inaccuracy of the raw sensor data into account. The spatio-temporal uncertainty of the fusion result is specified using uncertML, an XML schema developed by the INTAMAP (Interoperability and Automated Mapping) project to describe the statistics of uncertain data. As with the underlying sensors, the fusion procedure is described with the OGC sensor model language SensorML. In this way, the fusion procedure can be treated as a sensor, but with the important characteristic that its result is a coverage. The coverage can be visualized using a Map & Diagram service from the SANY partner ETH Zürich.

The procedures developed to date are variants of the Bayesian Maximum Entropy method that is able to consider soft sensor data (eg where the sensor

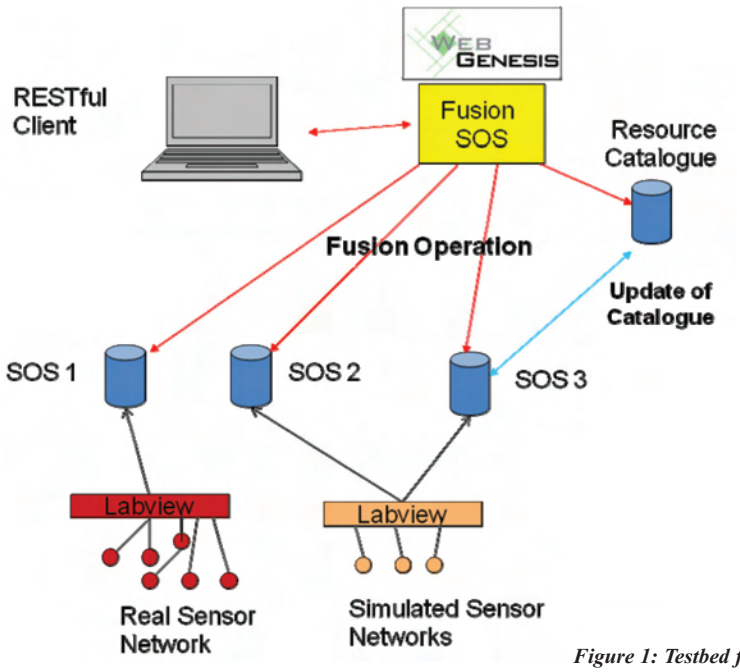


*Figure 1: Testbed for Sensor Service Networks.*

value lies in an interval) and additional phenomenological knowledge on the relationships between observed properties. If additional sensors or SOSs enter the testbed, the Fusion SOS discovers these new resources with the aid of the catalogue and incorporates the new data sources automatically into the fusion procedure. The self-describing information plays an essential role in this plug-and-measure capability.

The Fusion SOS is implemented on the platform WebGenesis, an information management server from Fraunhofer IITB. The information management server contains the information categories features of interest (sampling grids), procedures and results with associated metadata to support searching. The intermediate files produced by the fusion procedure are uploaded together with the fusion result to the WebGenesis information management server. This ensures a reproducible trace of the processing steps.

On-going work in the testbed involves the use of RESTful Web services to provide representations of the network resources and the development of model-based fusion methods. SANY (Sensors Anywhere) is an FP6 Integrated Project co-funded by the European Commission within the Thematic Priority 'Information Society Technologies' in the area of ICT for environmental risk management. SANY is a three-year project that started in September 2006 and has now completed two of three development cycles. The SANY consortium is composed of sixteen partners from eight countries. It includes the two research organizations Austrian Research Centers (coordinator of the consortium) and Fraunhofer, six companies, three universities, four public authorities and the Open Geospatial Consortium Europe (OGC).

**Links:**
http://www.sany-ip.eu/
http://www.opengeospatial.org/

**Please contact:**
Kym Watson
Fraunhofer Institute IITB, Germany
Tel: +49 721 6091 486
E-mail:
kym.watson@iitb.fraunhofer.de

# Bringing the Semantic Sensor Web to Smart Buildings

by Rob Brennan

***Deploying sensor networks in the built environment is not enough to produce smart buildings. How can we avoid creating silos of application-specific sensor networks? How can we publish the sensor data in secure, reusable and flexible ways? Can we support end-to-end provisioning of these sensor networks as an integral part of the building from requirements collection through design, procurement, construction, commissioning and facilities management operations?***

The Facilities Management (FM) subgroup of the NEMBES (Networked Embedded Systems) project aims to answer these questions and more by bringing together a multidisciplinary team of architects, civil engineers, computer scientists and electronic engineers, who are addressing the issue of networked embedded sensor systems in a holistic way across the stack from chip design, networking, middleware and service management. This four-year project started in October 2007 and is funded by the Irish Government's Higher Education Authority under the Program for Research in Third Level Institutions program. The project is led by the Centre for Adaptive Wireless Systems, and the FM subgroup is led by the Informatics Research Unit for Sustainable Engineering (IRUSE).

Traditionally, facilities management in the AEC (Architecture Engineering and Construction) domain is concerned with infrastructural facilities management (eg security and emergency management or space management) and technical facilities management (eg energy management, buildings operation and maintenance). Within NEMBES-FM we are extending the application of sensors in smart buildings with additional sensing capabilities and sensor-centric applications for these areas. However, rather than building application-specific sensor networks and ending up with isolated silos of non-interoperable and inflexible sensor networks, the project has adopted a semantic Sensor Web design to give a Web-based, open distributed system of sensor resources within the building. This enables resource sharing, resource reallocation, sensor network interoperability, sensor discovery and intelligent applications that discover and reason over associations, for example between events in space and time or within a particular context.

This is accomplished through the application of semantic Sensor Web technology such as SensorML (to describe sensor deployments and capabilities) developed by the OpenGS consortium, and standard W3C Semantic Web technology such as metadata definitions encoded as RDF (resource description framework) documents. This means adding formal semantic annotations to existing standard Sensor Web languages in order to provide semantic descriptions and enhanced access to sensor data. This is accomplished with model references to ontology concepts that provide more expressive descriptions of and relationships between concepts. The use of formal metadata to describe the sensors' outputs, platforms, locations and control parameters will enable a new generation of flexible facilities management applications to be built.

Efficient and flexible management of disparate, decentralized information sources such as sensor data, building occupancy graphs, facilities management process models and building information models will enable smart

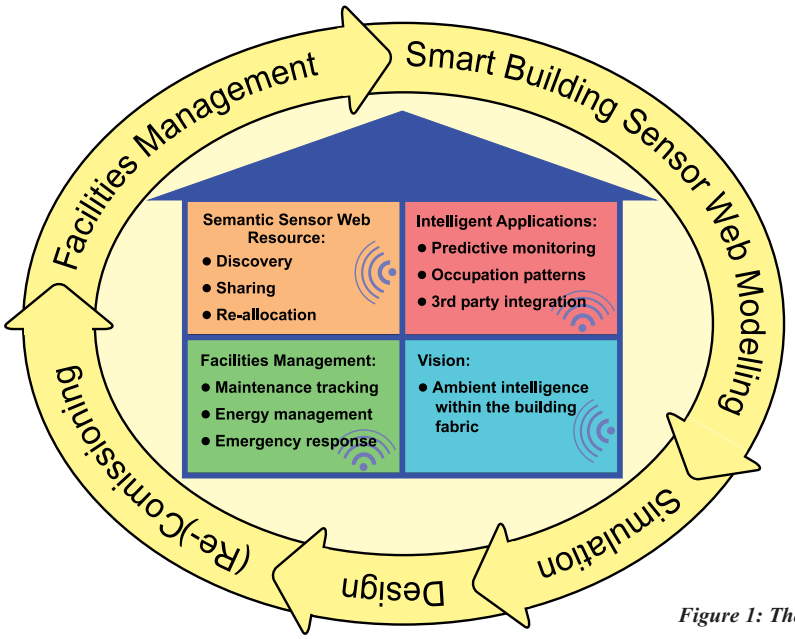


*Figure 1: The NEMBES project.*

FM applications for buildings with multiple occupying organizations and a distribution of facilities management authority across different management roles. The long-term vision of the group is to enable ambient intelligence within the smart building. In this scenario, building information models combined with location-sensing technology will allow the distribution of context-specific data to facilitate the monitoring of maintenance activity progress, ie the 'ambient interaction' of inspection and maintenance personnel with the fabric of the building itself. Examples of the benefits of this technology will include assisting facilities management staff with automatic monitoring and support of health and safety routine procedures; avoiding illegal occupation density in public spaces by triggering security personnel actions; and more effective facility space and relocation management via automatic inventory item tracking. Smart infrastructural FM applications that easily integrate new personnel, inventory items, sensors and use-cases without the intervention of dedicated IT personnel, will instead have these tasks performed by FM personnel and their delegates within the organizations occupying the building. They will additionally support seamless process integration between local FM activities and external third parties such as contract maintenance engineers, visitors, short-term occupants (eg conference organizers) and security or emergency response teams.

This project also goes beyond traditional ICT standardization to embrace the major AEC domain IT standards such as the IFC (Industry Foundation Classes), used for describing models of buildings in CAD tools and AEC construction requirements and project management tools. Integrating our work with the IFC standards will enable a dialogue with the ACE domain experts, thereby encouraging rapid integration of semantic Sensor Web technology in ACE tools and business practices. The project consortium is fortunate to include the Environmental Research Institute at UCC, a live smart building, in which sensor networks and building information models are combined with Semantic Web technology to provide a unique testbed.

**Links:**

The NEMBES project website:
http://www.nembes.org

The Knowledge and Data Engineering Group, Trinity College Dublin:
http://kdeg.cs.tcd.ie/

The Centre for Adaptive Wireless Systems, Cork Institute of Technology:
http://www.aws.cit.ie/

Informatics Research Unit for Sustainable Engineering (IRUSE)
http://zuse.ucc.ie/iruse/

**Please contact:**
Rob Brennan
Knowledge and Data Engineering Group, Trinity College Dublin, Ireland
Tel: +353 1 896 8426
E-mail: rob.brennan@cs.tcd.ie

# Building the Sensor Web – Standard by Standard

by Andrew Woolf

***An explosion in the instrumentation of our environment using sensors of all descriptions is driving the development of infrastructure to manage the wealth of information they collect. The Sensor Web aims to simplify the publication of, and access to, sensor resources, just as the World Wide Web has done for documents. And, as with the WWW, the Sensor Web relies on new information and communication standards for structuring sensor information and its exchange.***

At the heart of these new standards is a conceptual model for Observations and Measurements (O&M, [OGC 07-022r1, OGC 07-002r3]). It says simply that an Observation is an action whose Result is an estimate of the value of some Property of a Feature-of-interest, obtained using a specified Procedure (Figure 1: The 'Observations and Measurements' conceptual model.). Each of these core O&M information classes may be extended to specific sensor applications (see box).

The abstract O&M model may be applied across the spectrum of sensor applications and deployments, and provides a framework for building exchange standards and service interfaces for accessing sensor data and contextual information (Figure 2: The Sensor Web standards stack.). For example, the Sensor Observation Service (SOS, [OGC 06-009r6]) provides a Web service interface for retrieving filtered observations or related information (feature-of-interest, sensor parameters, observation results). Individual sensor observations may be aggregated within one service into combined 'observation offerings' and multiple services may be federated into single access points. The Sensor Model Language (SensorML, [OGC 07-000]) is an XML language for describing observation procedures and sensor types. Other related standards include the Transducer Markup Language (TML, [OGC 06-010r6]) for transducers and transducer systems, the Sensor Planning Service (SPS, [OGC 07-014r3]) for tasking and scheduling observation requests with sensor systems (eg by satellite remote-sensing instruments), and the Sensor Alert Service (SAS, [OGC 06-028r3]) for setting up notification subscriptions for specific sensor events.

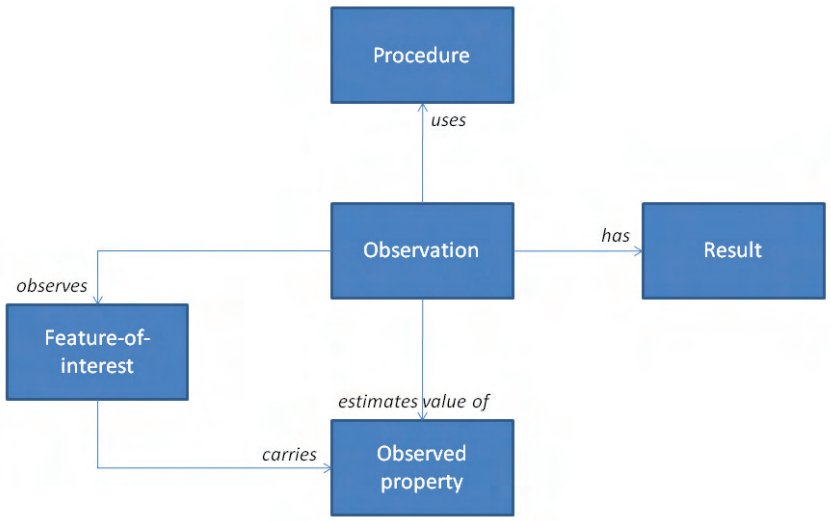


*Figure 1: The 'Observations and Measurements' conceptual model.*

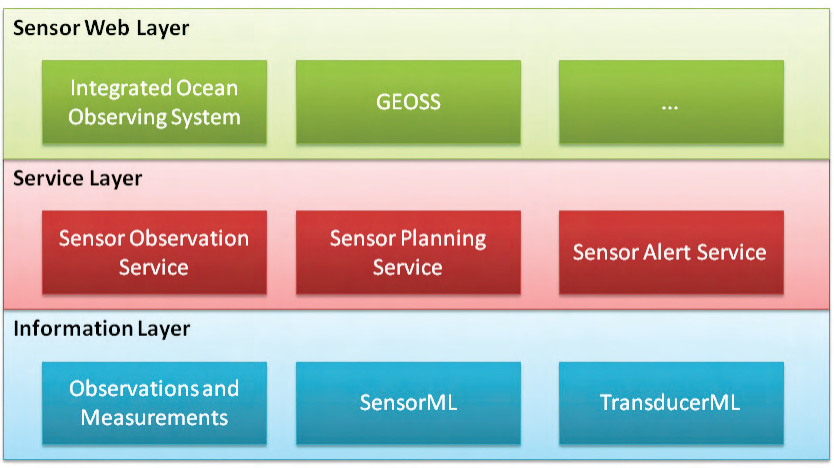


*Figure 2: Figure Two: "The Sensor Web standards stack.*

These new standards are being developed by the Open Geospatial Consortium (OGC), a non-profit de facto international standards body for geographic information. OGC members span government, academia and industry. Organizations contributing to the O&M specification include CSIRO (AU) as editor, Geoscience Australia (AU), University of Alabama in Huntsville (US), Image Matters LLC (US), Washington University (US), Science and Technology Facilities Council (UK), SeiCorp Inc. (US), Galdos Systems Inc. (CA), Geospatial Research & Consulting (DE), PCI Geomatics (CA), and Texas A&M University (US).

Research undertaken through global collaborations of experts, often in the context of coordinated international engineering testbed activities (the OGC Web Service Initiatives) accelerates the development of interoperability stan-

dards and protocols. With sufficient maturity, these specifications will be proposed for de jure standardization through the International Organization for Standardization (ISO).

The O&M conceptual model has been substantially developed since 2002 through such initiatives (OWS-1.2 and OWS-3). It was finally approved in 2008 for Version One publication by OGC, and is being progressed as a new ISO standard 19156 (see links below).

A growing awareness of the importance of environmental monitoring for the health of our planet is leading to the development of large-scale infrastructures that transcend national boundaries. Through a ten-year implementation plan, the Group on Earth Observations (a grouping of 76 national governments and other international organizations) aims to integrate existing observation networks into a Global Earth Observation System of Systems (GEOSS) to achieve comprehensive, coordinated and sustained observation of the Earth system.

Within Europe, a recent Directive (2007/2/EC) will establish the 'Infrastructure for Spatial Information in Europe' (INSPIRE) to integrate environmental data across all member states. The 'Kopernikus' partnership between the European Commission and the European Space Agency will establish core operational services (eg ocean forecasting, landcover monitoring, emergency response) for the global environment and civil security. These three global-scale initiatives all require standard information models and network services for integrating sensor data – both in situ and remotely sensed (spaceborne and airborne).

An underlying abstract architectural approach is used to develop Sensor Web standards. The Reference Model for Open Distributed Processing (RM-ODP, [ISO/IEC 10746]) factors a distributed system like a Sensor Web into five complementary viewpoints: enterprise (roles, scope and policies of the system), information (semantics of information and information processing), computational (service interfaces), engineering (component distribution across nodes) and technology (implementation choices). Against this model, O&M provides the information viewpoint, while service standards like SOS and SPS provide the computational viewpoint.

Within the broader context of ICT innovation, Sensor Web standards build on key W3C specifications (XML, Web services). There is a growing move towards richer semantics, for describing observed properties and their relationships; for instance, ISO 19150 is a new standard for the use of ontologies with geographic information. An important principle is the decomposition of domains of governance – identifying responsible parties with a remit for managing agreed vocabularies and concepts on behalf of a community of interest (eg definitions of particular sensor systems and observables). This is required in order to facilitate reuse of models and enhance interoperability.

While much of the Sensor Web standardization work has so far been developmental, there are beginning to emerge practical demonstrations of their effectiveness. Demonstrators like the US-based OpenIOOS testbed for an Integrated Ocean Observing System and projects like the European 'Sensors Anywhere' (SANY) FP6 Integrated Project are validating the application of these standards in real Sensor Web applications.

## Example: Marine Science Observation

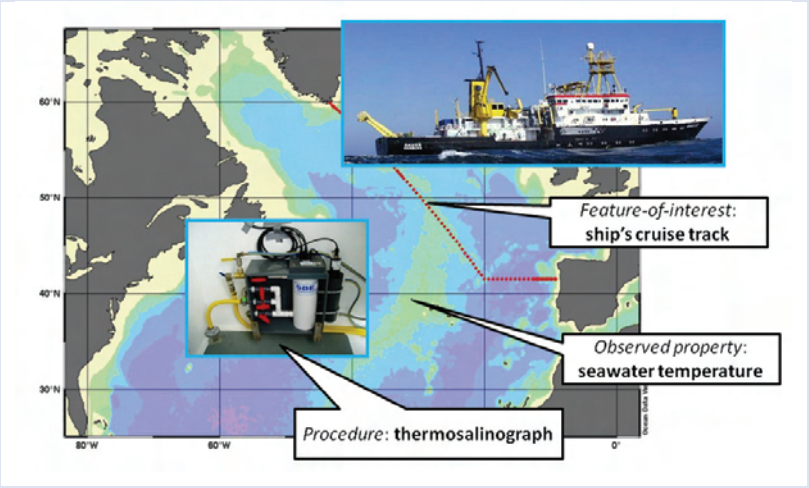


***Figure 3: O&M example – marine research cruise.***

A cruise ship measures seawater temperature (O&M observed Property) along a North Atlantic cruise track (O&M Feature-of-interest) using a thermosalinograph (O&M Procedure), and produces a series of data values {19.2°C, 18.7°C, ...} (O&M Observation result).

The SOS *GetObservation* operation is used to retrieve the entire observation and its context; *DescribeSensor* will return thermosalinograph details; *GetFeatureOfInterest* provides details of the cruise track; while *GetResult* returns just numerical temperature values.

**Links:**
http://www.opengeospatial.org/
http://www.earthobservations.org
http://www.opengeospatial.org/ogc/markets-technologies/swe
http://www.isotc211.org/
http://inspire.jrc.ec.europa.eu/
http://ec.europa.eu/gmes/index_en.htm
http://www.sany-ip.eu

**Please contact:**
Andrew Woolf
e-Science Centre, STFC Rutherford Appleton Laboratory, UK
Tel: +44 1235 778027
E-mail: Andrew.Woolf@stfc.ac.uk

# Revolutionising Sensor Based Automation in Manufacturing

by R Harrison, F Jammes, H Smit and T Kirkham

***Increased access to device-level automation components is closing the final gap in the enterprise computing model.***

The implementation of Web service-enabled sensors and actuators on production lines will permanently change the way in which future automation systems are designed and implemented. Current interfaces to automation components are largely vendor-specific, restricting the reconfiguration of lines and the management of line data across enterprises. Outside the automation domain, enterprise system development has seen real-time data linkage take great steps in the office, warehouse and supply chain. Research into the use of Web-service-enabled sensors and actuators has the potential to present an open standards-based method to integrate production lines into this enterprise computing model, an innovation that will revolutionize automation in future manufacturing plants and remove control from vendors back to users.

## Implementation

The SOCRADES and SODA (Service-Oriented Device and Delivery Architecture) projects have conducted research and trials into automation based on service-oriented architecture (SOA), and in the past year have delivered initial prototypes. At the recent ITEA exhibition in Rotterdam, a test rig developed with Ford was used to demonstrate SOA-based sensor and actuator data being used to manage monitoring and control applications. The data was fed to project partners in the Enterprise computing field (SAP) and industrial automation sector (ARC Informatique). The demonstrator illustrated how the production data could be combined with other enterprise data to improve the accuracy of decisions relating to production routing and supply chain management.

During the demonstration, data from the sensors and actuators on the production line was transmitted by equipment created by Schneider Electric. This equipment consisted of Web-service-enabled Field Terminal Blocks (FTBs), which support the Device Profile for Web Services (DPWS) toolkit. The DPWS toolkit is designed for embedded systems and has a small memory footprint, but also contains a selection of Web service standards to suit the demands of an automation environment. Both projects are working on

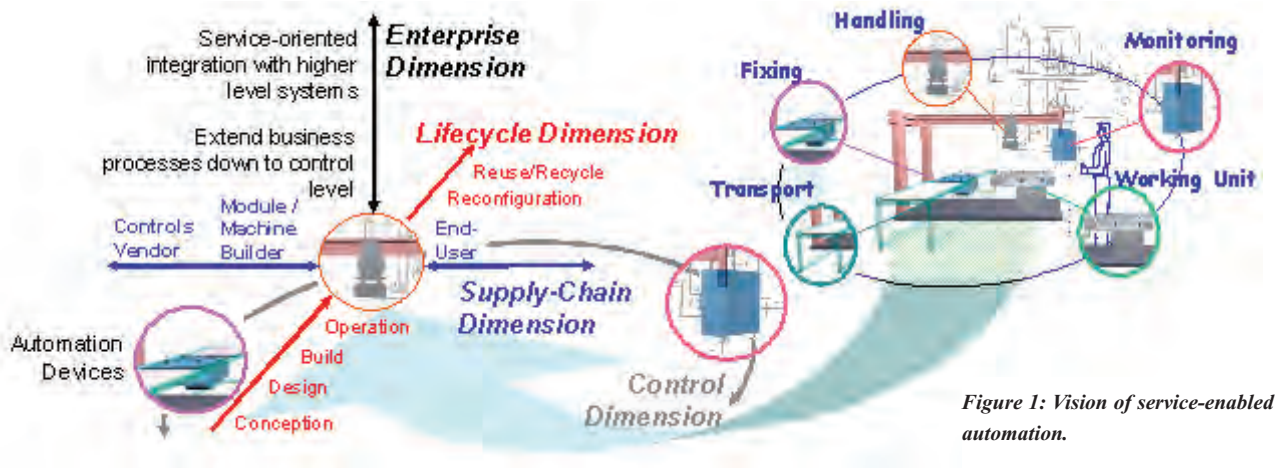


*Figure 1: Vision of service-enabled automation.*

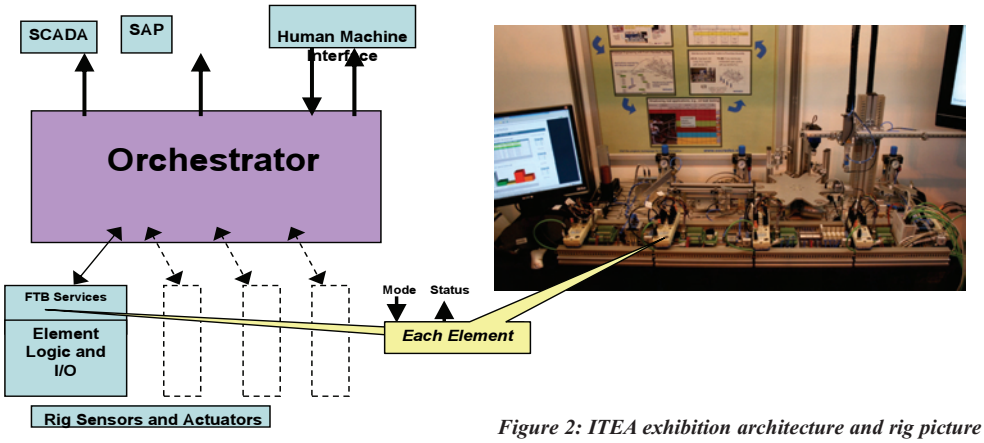


*Figure 2: ITEA exhibition architecture and rig picture showing FTB location.*

these demands, which are focused on execution timing and reliable/efficient message delivery. The FTB is a piece of hardware that contains an ARM 9 chip developed to support the DPWS toolkit.

The support of DPWS on the FTB allows Web service interfaces to be developed to the device-level I/O within the line. For example, calls to directly command an actuator or monitor specific sensors on lines can be made via Web services located on a variety of applications, as opposed to specific vendor control devices. The control of devices by Web services has been achieved using central Web service orchestration and also on a smaller-scale peer-to-peer choreography.

### Future

Device-level automation components that produce data in standard and open forms will be faster to reconfigure, reducing costs by improving resource usage and reducing downtime. The greater accuracy and real-time access to product-level data will further enhance enterprises by allowing them to make more accurate decisions regarding production and supply chain matters. The application of the results from these research projects in real manufacturing environments will be the subject of future work that should confirm these findings. The adoption of the approach will also be dependent on further research in the areas of safety, security and real-time execution of devices.

Direct SOA linkage to sensors and actuators moves a traditionally vendor-specific computing area into a new open domain, ready to link with existing innovations in enterprise computing. For the manufacturer this will improve performance and reduce costs. However, this is dependent on the continued development of devices such as the FTB, pioneered in the SOCRADES and SODA projects to support this new vision.

**Links:**

SOCRADES: http://www.socrades.eu

SODA: http://www.soda-itea.org

Device Profile for Web Services home page including link to specification: http://schemas.xmlsoap.org/ws/2006/02/devprof/

**Please contact:**
Tom Kirkham
Loughborough University,
Leicestershire, UK
E-mail: T.D.Kirkham@lboro.ac.uk

# Security and Robustness of Wireless Sensor Networks

by Václav Matyáš and Petr Švenda

***Researchers at Masaryk University, Brno, are working on security issues relating to large-scale, highly distributed and relatively dense wireless sensor networks.***

In our work we focus on link key establishment in the memory- and computation-restricted environment of wireless sensor networks (WSNs). We also study how link security behaves under a selected attack and what methods can be used to strengthen the resilience of WSNs against compromise. We base our work on the assumption that a partial compromise in WSNs is inevitable and network architecture should be prepared to cope with related security issues. We work with two basic link key establishment concepts based on symmetric cryptography: memory-efficient probabilistic pre-distributions (Eschenauer & Gligor, 2002) and lightweight key exchange without pre-distributed secrets (Anderson et al, 2004). These two key distribution concepts behave differently when the network is attacked. Analysis of the resulting compromised patterns has led to the proposal of mechanisms for improving the network resiliency based on support from neighbouring nodes.

While the resiliency of probabilistic pre-distribution schemes generally increases when more keys can be put into a key ring on every single node, such an increase is limited by the node storage capacity. Our multiparty protocol creates a large virtual key ring in an efficient and secure way from the key rings of separate nodes. This results in a substantial increase in resilience of the underlying probabilistic key pre-distribution scheme against the threat of node capturing. The protocol performs similarly to the hypercube pre-distribution (Liu & Ning, 2003) but is more suitable for scenarios with random deployment and unknown link compromise status. The proposed protocol itself is also resilient against partial compromise inside a group of supporting neighbours.

Our former work exploited non-uniformity of link compromise patterns in key infection, and led to a secrecy amplification (SA) protocol with a significantly better fraction of secure links than previously published SA protocols, especially for denser networks. We applied SA protocols of partially compromised networks resulting from node capture when probabilistic key pre-distribution are used, and provided analytical and simulation evidence that SA protocols work even better here. On average, SA protocols secure more links for probabilistic pre-distribution than for key infection, when networks with the same percentage of initially compromised links are assumed. When the SA protocols are applied, a network with half of its links compromised can be made reasonably secure with less than 10% of compromised links.

Some combinations of SA protocols that worked for key infection do not increase the number of secure links in probabilistic pre-distribution and thus only impose unnecessary communications overhead. Instead of analysing

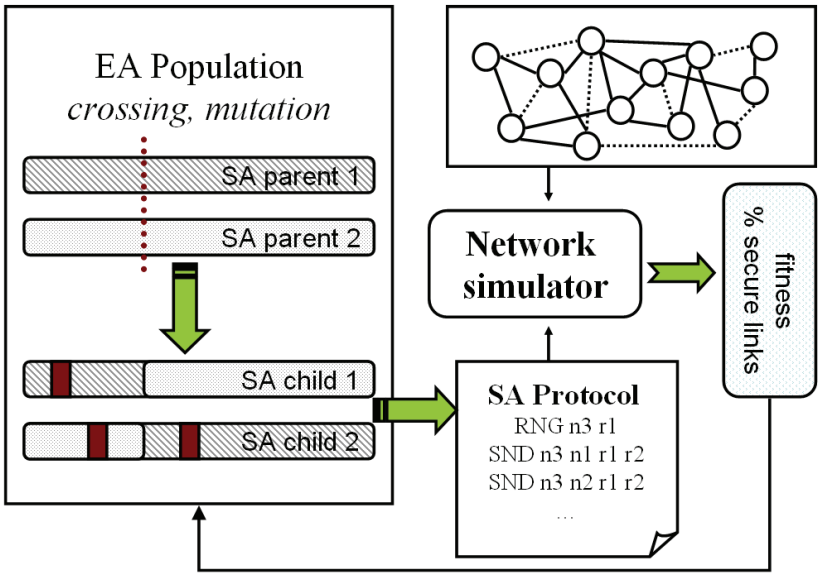


*Figure 1: Automatic generation of secrecy amplification protocols.*

each separate compromise pattern arising from the combination of a particular key distribution method and attacker strategy, we proposed an automated approach based on the combination of a protocol generator and network simulator. We utilize evolutionary algorithms to facilitate guided searches for high-performance SA protocols created as a series of elementary instructions. Every candidate protocol is evaluated on our network simulator for a particular compromise pattern (see Figure 1).

Using this method, we were able to automatically re-invent all the human-designed SA protocols of which we were aware, and to find a new protocol that outperforms these. Moreover, we proposed an alternative construction of SA protocols that exhibits only a linear (instead of exponential) increase in necessary messages when the number of neighbours in the communication range (network density) is growing, and we achieved comparable performance to protocols with original message-expensive assumptions providing energy-efficient SA protocols. With respect to classical human-made protocols, an increase in the number of secure links was obtained by an efficient combination of the simpler protocols and an unconventional interleaving of elementary instructions. These allow protocols to be executed even when one of the participants is out of radio transmission range.

Our current work focuses on the concept of automatic search for attack strategies with demonstrative applications to link key security for probabilistic pre-distribution and key infection approach. New attacks are generated either as a recombination of existing attacks or as completely novel attacks automatically assembled from elementary attacker actions. They are then evaluated on a network simulator or in a real system. Attacker strategies that increase the number of compromised links with respect to several deterministic algorithms or random cases were found. Initial results for attacks against selected routing protocols show good prospects for an automated search for selective jamming, message dropping and neighbour overloading to achieve a specific attacker goal such as increasing routing path length, message latency or concentrating routed messages.

Due to battery power limits and taking into consideration the high communications overhead exposed by current replication detection, the reputation management mechanisms that have been proposed so far are often not affordable. We are currently designing prevention, detection and reaction techniques for the network. Rather than aiming for perfect security, which is particularly hard to achieve in WSNs, the aim is to force attackers to make disadvantageous trade-offs in terms of computational time, energy or other costs. Due to the diversity of usage scenarios, there is a need to develop an economical/mathematical model that would help to find a near-optimal solution for a particular combination of network usage and available resources.

**Link:**
A technical report covering some of the issues discussed above:
http://www.fi.muni.cz/reports/files/2007/FIMU-RS-2007-05.pdf

**Please contact:**
Václav Matyáš
Masaryk University, Brno – CRCIM, Czech Republic
Tel: +420 549 49 5165
E-mail: matyas@fi.muni.cz

# Security Challenges for Wireless Sensor Networks – Dynamic Routing as a Security Paradigm

by Marek Klonowski, Michał Koza and Mirosław Kutyłowski

***Recent work carried out at Wroclaw University of Technology shows that a fair level of security can be achieved for wireless sensor networks without heavy cryptographic technology.***

Wireless sensor networks processing sensitive data are facing the risks of data manipulation, data fraud and sensor destruction or replacement. This concerns applications such as the gathering of data on environmental pollution around industrial installations, or sensor systems replacing traditional video monitoring. Large-scale deployment in practice is conditioned by solving these kinds of security problem and reducing the risks due to limited physical protection of the devices and openness of the wireless communication channel. While modern cryptography and computer security offer many ways of solving these problems, they are focused on solutions for high-performance devices, and not for computationally weak sensors with limited communication bandwidth. New 'lightweight' solutions tailored for the special needs of wireless sensor networks have to be designed. This is one of the focal points of the EU project FRONTS (Foundations of Adaptive Networked Societies of Tiny Artefacts). Fortunately, some recent developments have shown that without heavy cryptographic technology it is still possible to achieve a fair level of security in a practical sense. This report indicates a few ideas of this kind.

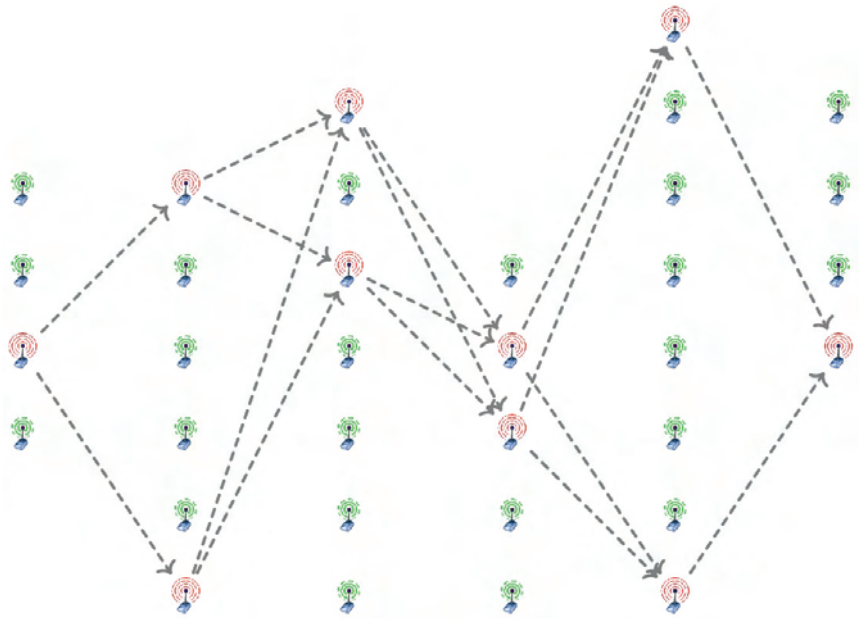

*Figure 1: Snapshot of a single routing path.*

Due to the energy required for transmission over long distances, it is often a good idea to route data along a sensor network by making many hops over small distances instead of a direct transmission from a sensor to the sink node. However, such a solution has the disadvantage that an adversary can attack the network by gaining control over intermediate sensor nodes. The cryptography used by such devices is usually weak and can provide opportunities to reveal information sent or to manipulate them.

The following idea may be applied in order to make it much more difficult to carry out attacks. Instead of a single information path, each message is sent over a double path. This means that instead of a single $i^{\text{th}}$ node $N_i$ we have two nodes: $P_i$ and $R_i$. The encryption scheme has the following basic properties when processing a message $M$:

- $P_{i+1}$ receives encrypted messages from $P_i$ and $R_i$ in order to compute its share of message M,
- $R_{i+1}$ receives different encrypted messages from $P_i$ and $R_i$ in order to compute its share of $M$.

The encryption scheme guarantees that corrupting either $P_i$ or $R_i$ reveals no information about $M$. Also, combining the shares from different stages of message processing gives no information

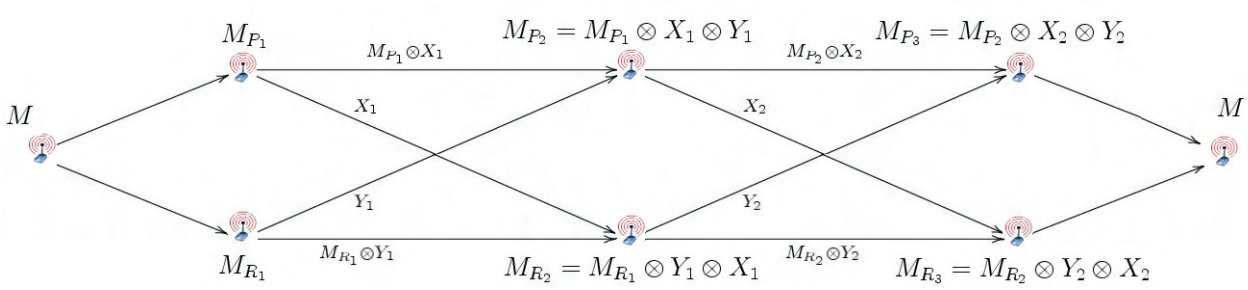


*Figure 2: Encoding details.*

about $M$ as long as the adversary has only one share from each level of the path.

What is the advantage of such a design? The main point is that while it might be relatively easy to find and corrupt one of the nodes (say $P_i$ ) for this to be useful, the adversary must still find and corrupt the matching node $R_i$ . This can be difficult for purely practical reasons: if each sensor is hidden in the environment, then while the first might be found by chance, the second must be found by a detailed search in the same area. This could be hard without arousing the interest of observers.

Moreover, we propose a far more sophisticated design in which on each level of routing there are many potential sensors to play the roles of $P_i$ and $R_i$. In this case the adversary usually has to collect and corrupt many sensors until the matching pair is found.

Another step of the design is to make the path self-evolving: at any time a node may negotiate with its predecessors and successors a change of the transmission key and redirect its duties to another node. Since these changes can be made independently and uniformly at random, the data path may evolve so fast as to make unfeasible any attempt at data analysis based on monitoring radio traffic. Indeed, a cryptanalytic attack would face the difficulty that assigning the messages to sensor-to-sensor links (and to the pairwise keys) would be hard, due to the number of possibilities growing extremely fast as the number of links increases.

The architecture described here is currently being analysed from the point of view of hiding the transmission routes in the case of heavy traffic, under the assumption that an adversary can select the traffic coming out of each node. This involves studies concerning combinatorial issues of traffic analysis as well as stochastic investigations of the rapid mixing of Markov chains. Further details of the scheme will be developed in cooperation with other partners of the project; in particular, we plan to develop a prototype of the system.

**Link:**
FRONTS: http://fronts.cti.gr/

**Please contact:**
Mirosław Kutyłowski
Wrocław University of Technology, Poland
Tel: +48 71 320 21 09
E-mail:
miroslaw.kutylowski@pwr.wroc.pl

# Ubiquitous Machine-to-Machine Service Networks

by Johanna Kallio and Juhani Latvakoski

***In the near future, there will be many more embedded devices than there are mobile phones. When these devices are connected to the Internet, many novel kinds of ubiquitous service will be enabled.***

It has been estimated that in 2010, the number of communicating devices will be a thousand times greater than the number of mobile phones, which is already more than one billion. When connecting devices such as various machines, actuators and sensors to the Internet, novel types of service are enabled. Previously, such devices communicated with services using technology such as SMS. The applications were vendor or domain-specific closed systems, for which achieving interoperability with other vendor/domain systems was challenging. The Usenet (Ubiquitous M2M Service Networks) project aims to enable ubiquitous machine-to-machine (M2M) service networks, in which the M2M infrastructure is able to connect and combine services produced in different domains in an interoperable way (see Figure 1).

## The Usenet Project

Currently, no universally applicable M2M service infrastructure exists that would allow interoperation between devices and their enabled applications in wired and wireless systems, regardless of the supplier. Information technology applications usually operate as separate M2M solutions that are unaware of each other. As a result, a number of business opportunities remain unexploited as the services provided by the devices cannot be placed on the Internet.

The three-year Usenet project funded by the Eureka/ITEA2 programme is developing a service concept for solving the above interoperability problems. M2M services refer to the services resulting from collection, transmission and processing of information, and establish an interactive system with the remote devices that are ultimately integrated within a managed M2M software system. The project has generated new types of M2M service scenario, which are related to ubiquitous building infrastructure, machine tools, consumer devices, home automation and telematics domains. The primary goal is to specify a universally applicable M2M concept that will enable the interoperation of sophisticated M2M applications through heterogeneous wired and wireless IP communication networks.

Made up of seventeen partners, the international Usenet consortium focuses on M2M research enabling ubiquitous M2M service networks. The project consortium is led by the VTT and includes industrial, SME and research partners from Finland, Belgium, France and Spain.

## M2M Architecture

The system components provided by different suppliers have a strong influence on the structure of M2M systems. M2M systems usually require the integration of components coming from various stakeholders in the value chain: M2M service providers, M2M operators, M2M manufacturers, software houses and M2M system integrators. The referred components need to be interoperable in order to establish sensi-

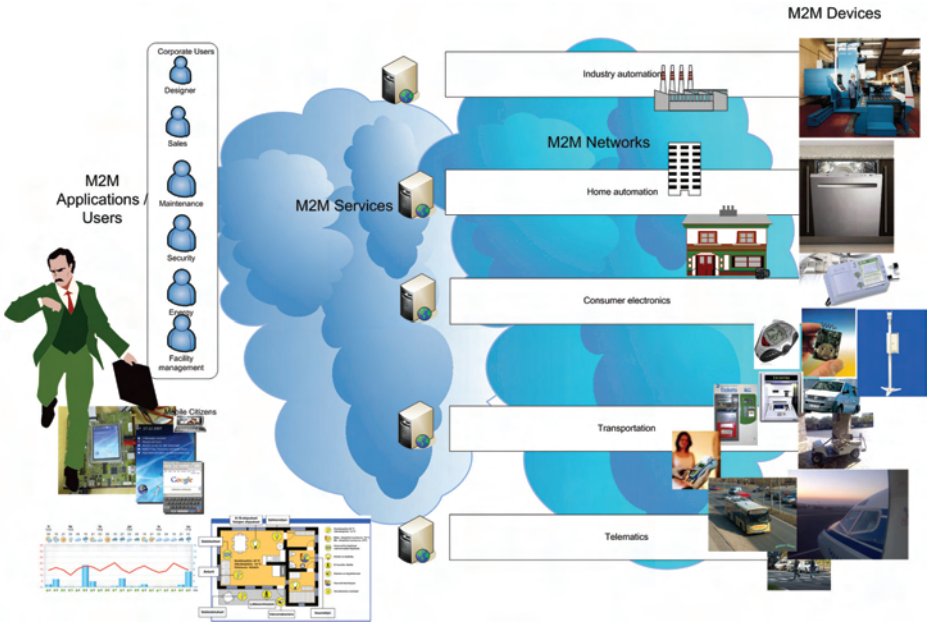


*Figure 1: Ubiquitous M2M service networks.*

ble business operations. Traditionally, M2M solutions have applied vertical architecture and closed solutions. This has created challenges in the distribution of added value, which has been a barrier to M2M market enlargement. It has been estimated that horizontal architecture has better possibilities to boost M2M market.

It is expected that horizontal architecture will make it easier for different players to be part of the M2M value network. For example, an M2M asset devices manufacturer can offer control and administration services for their products. Communication infrastructure can be connected to these devices by means of various telecommunications' manufacturers and service providers. Service platforms can utilize several communication infrastructures for collecting data and controlling M2M devices. Smart services can be based on information, which is collected from several service platforms. The challenges of horizontal systems are related; e.g. the overall quality of end user services and security, which inevitably requires the existence of vertical interfaces.

## Home Surveillance – an M2M Application Example

A Usenet M2M application scenario has been implemented to demonstrate the Private Space M2M system for residential homes. The purpose of the system is to provide smart surveillance services of private space. Various sensors measure quantities such as humidity, temperature, light levels and consumption of water or electricity, and the space is monitored with a video camera. The user is able to follow what is happening in the system via the home user interface (Figure 2), and can control the space with sensors and actuators connected to the space. For example, users can track the outside and inside temperatures and are provided with alarms and warnings of water leaks, unexpected weather or the presence of a housebreaker. They also have control over the lights, heating and so on. Automatic services related to control operations can also be included.

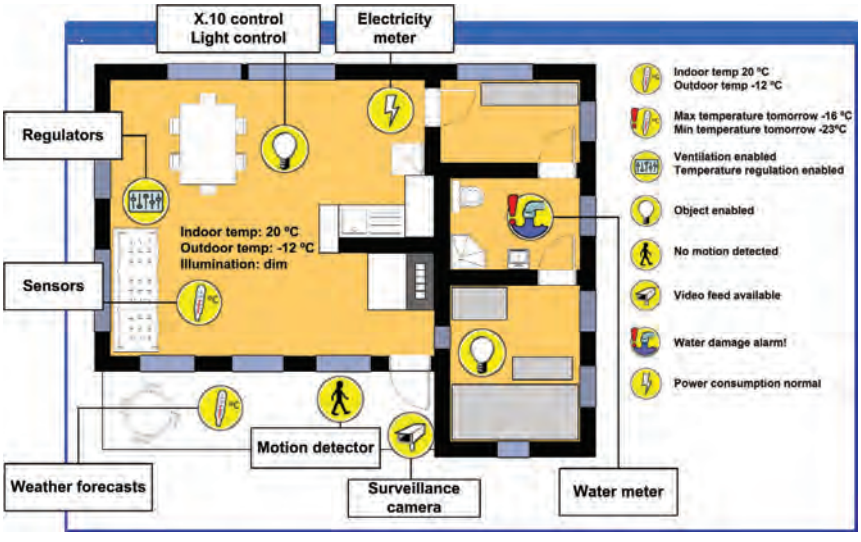


*Figure 2: User interface of Usenet experimental M2M application.*

## Potential of M2M Systems

M2M systems will provide essential business possibilities and advantages for companies, especially when information systems controlling their core processes are utilizing the real-time information produced by an M2M system. In consequence, a company can increase the quality of its services, reduce costs and increase customer satisfaction. This fundamental change, which will bring new business opportunities for companies, can already be seen in the market. VTT aims to help companies to take advantage of this rapidly growing M2M market.

**Link:**
Usenet project: https://usenet.erve.vtt.fi/

**Please contact:**
Juhani Latvakoski, Johanna Kallio
VTT, Finland
Tel: +358 40 520 0149
E-mail: Juhani.Latvakoski@vtt.fi, Johanna.Kallio@vtt.fi

# Short-Range Communication in Ubiquitous Professional and Consumer Applications

by Mikko Sallinen, Esko Strömmer and Pirkka Tukeva

***Short-range communication technology is proving its worth in many areas of application. Here we illustrate three case studies involving NFC (Near-Field Communication) and Nanonet radio technology. These include a consumer application with the TouchMe paradigm, and professional applications for force measurements on train wheels and acceleration measurements on the axis of a paper mill machine. We also show the advantages of these short-range communication technologies and discuss other potential applications.***

In recent years, a rapidly growing range of wireless communication technology has become available and is being applied to numerous application domains. One of the latest examples of short-range communication technology is a light and simplified version of Bluetooth known as Bluetooth ULP (Ultra-Low Power), earlier known as Wibree. Other examples of short-range communication technology include RFID, Bluetooth, IrDA and Zigbee. From these technologies, IrDA and Bluetooth offer a huge number of applications.

New technology will also generate a new business for the service sector which is an important growing market. There are many areas in which wireless technology will find ubiquitous application, including entertainment, healthcare, automotive and logistics. These are areas which have a huge future market potential.

In industrial or other professional applications, the requirements for accuracy, reliability and timing are typically higher than in consumer applications. Current applications include measurement and wireless data transfer from the field to a central unit, as well as control and monitoring applications. In many cases this works well, but the problem is often that even if the data can be transferred, it is still necessary to have power cables. Energy harvesting and power management is one of the key issues in the more challenging industrial applications.

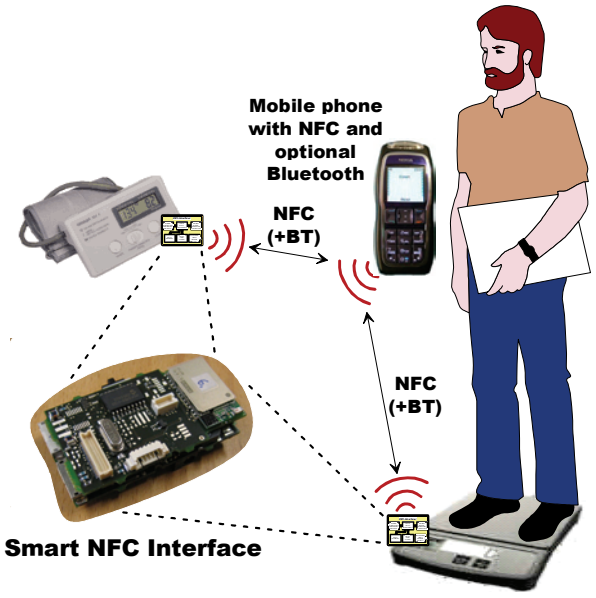


*Figure 1: Intelligent interfacing module for smart objects.*

In this article, we use two different technologies. The first is NFC (Near-Field Communication), an RFID-based technology that provides short-range communication at 13.56 MHz licence-free band. The second is the Nanonet radio module, which uses the 2.4 GHz licence-free band and is manufactured by Nanotron Ltd.

## Touch-based Sensor Readout by NFC

NFC can be used to read data from a sensor to a mobile handset by touching the sensor to the mobile handset. In more data-intensive sensor applications, NFC can easily establish a Bluetooth connection between two devices when they are brought close together. It is not necessary for a user to find the correct menu items and configuration parameters on the mobile handset, which makes interaction with the NFC-enabled sensor very elegant. On the other hand, the sensor must be within the reach of the user.

For easy piloting of NFC applications we have developed Smart NFC Interface, a building block that provides NFC and Bluetooth communication capability to various smart objects. Thus it also makes these smart objects capable of communicating with mobile handsets as depicted in Figure 1. Smart NFC Interface incorporates components for short-range wireless communication with mobile handsets, a set of wired interfaces for sensors and other smart objects, local data processing and storage, and a clock for time stamping the measurement results and other events.

The advantages of NFC over alternative wireless short-range communication technologies in sensor applications, such as Bluetooth, are its lower power consumption, shorter communication set-up latency and better immunity to eavesdropping. NFC can also operate in passive mode with extremely low power consumption, which means wireless sensors can have multiple-year battery lifetimes. Even batteryless sensors taking their power from the RF field that is generated by a nearby active NFC device are feasible.

In addition, NFC enables simple communication based on the TouchMe paradigm that from the user's viewpoint is fast and convenient compared to plugging in cables or manually establishing a Bluetooth connection between the sensor and the mobile handset.

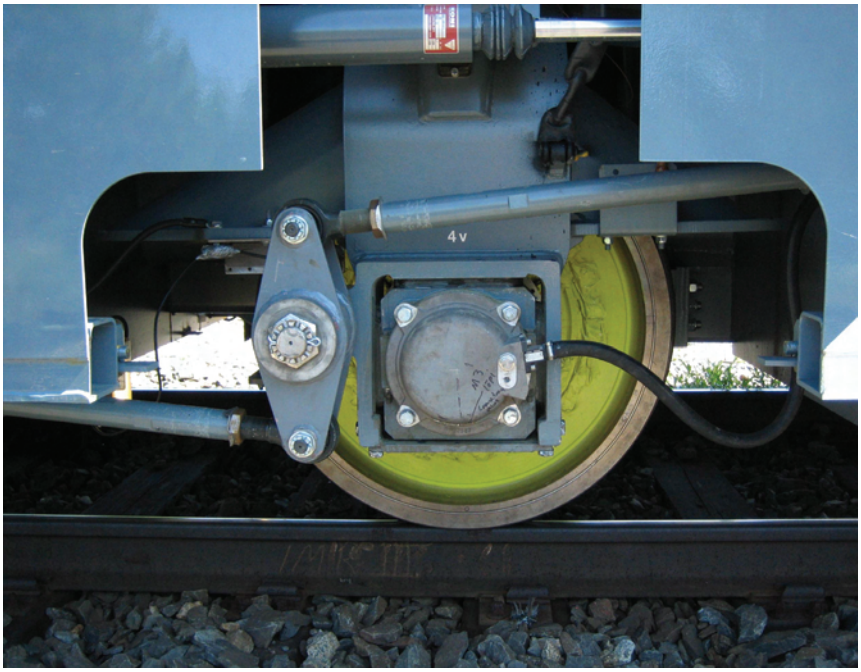

*Figure 2. Force measurements from a train wheel.*

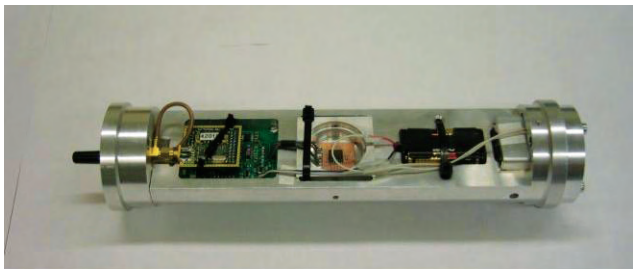

*Figure 3: Measurement system inside the paper roll.*

## Force Measurements on Train Wheels

A typical challenge in the industry is taking measurements on rotating machines. The difficulty of wiring both data and power supply makes this a non-trivial problem. Examples of this kind of object are wheels, turbines, blades and paper rolls. In this field, the goals are to optimize construction and maintenance by condition monitoring, and to be able to foresee the need for maintenance. It would be much more effective to carry out maintenance based on the actual wear of parts, rather than based on age as tends to be the case. In the future, this will lead to significant savings.

To take force measurements on train wheels, strain gauges are placed on the surface of the wheel in the form of a star. The measurements are converted using FPGA platform and then transferred using the 2.4 GHz ISM band and 2Mb/s wireless link. The power consumption while in use is 60 mA, and 1µA while in standby. Maximum transit power is 10 mW. The outdoor range of the radio is 100m; when indoors it depends on the presence of surfaces and walls. Communication with sensors is carried out using an SPI (serial peripheral interface) interface with a 16 Mhz line. This wireless link is carried out using NanoNET power modules by Nanotron.

Nanotron uses 'chirp' modulation technology, which provides benefits such as resistance to multipath fading. In the receiver side of the measurement system, the System-on-Chip platform receives the data. In each wheel there are twenty pairs of strain gauges that measure the deformation of the wheel. These strain gauges are read at a frequency of 3 kHz. The measurements are filtered to a frequency of 500 Hz and the data is transferred using 48 channels for transformation. It is therefore necessary to run several radios in parallel and to carry out high-speed synchronization. The sampling rate is 290 MHz in the developed system.

## Acceleration Measurements from the Axis of a Paper Mill

In the third application, we integrated a 3D acceleration measurement system with the axis of a paper roll. The goal was to measure the acceleration of the axis and be able to predict cases in which the paper will tear in the paper machine. It is known that tear will occur when there is vibration in the axis.
For this purpose, a specific measurement tool was designed (see Figure 3). It includes a triple-axis acceleration sensor and a transmitter to send online measurement signals. The challenge in the task was to successfully transmit a signal through a full paper roll. For the data transmission, we used the same radio technology as in the train wheel case; with this we were able to measure and recognize peaks in the vibrations of the axis and the data was transmitted successfully even through the thickness of the roll.

**Please contact:**
Mikko Sallinen
VTT Technical Research Centre, Finland
Tel: +358 40 7235263
E-mail: Mikko.Sallinen@vtt.fi

# The CLARITY Ubiquitous Robotic Testbed

by Gregory O'Hare, Mauro Dragone and Jennifer Treanor

***The CLARITY Centre for Sensor Web Technology in Ireland is currently constructing a ubiquitous robotics testbed by integrating a collective of mobile robots with a wireless sensor network and a number of portable devices. The new, mixed testbed will be hosted at the School of Computer Science and Informatics at University College Dublin, (UCD), and will also avail itself of the laboratory facilities hosted in Dublin City University (DCU) and Tyndall, Cork. The testbed will provide a service for all researchers interested in developing ubiquitous robot applications.***

The number and variety of applications of robots in our daily environment is on the increase. Examples include robotic pets (eg Sony Aibo), household appliances (eg iRobot vacuum cleaner Roomba) and assistive technology (eg the MANUS wheelchair). At the end of 2006, figures for service robots for personal or domestic use stood at nearly 3.5 million, with this projected to more than double before the end of 2010.

At this point, a principal challenge for the robotics community is the integration of robots into today's digital society. Inexpensive Internet access and the diffusion of wireless computing devices have made ubiquitous/pervasive computing a viable reality that augments the normal physical environment and supports the delivery of services anytime, anywhere to human users. Endowing these ubiquitous devices with intelligent behaviour, and thus creating intelligent environments, is termed ambient intelligence.

Robots are a compelling instance of those artefacts that comprise and deliver the ambient space. Modern multi-robot applications have moved away from the historical view of monolithic control systems: they run on specific computational units and are in charge of their own 'hardwired' hardware. They are already perceived to belong to a larger, open distributed network made up of different sensors and effectors. Ubiquitous robotics further extends this view by explicitly addressing the need for interoperability between robots and existing ubiquitous and pervasive infrastructures, such as wireless sensor networks (WSNs).

CLARITY is particularly interested in employing the new testbed in three overlapping areas of research:

## Integration between Robots and WSNs

The new testbed represents an ideal arena in which to develop and test proposed advances in mobile networking, routing, data collection and data analysis within the WSN community. Robots are just another class of user of sensor data, potentially harvested both by their on-board sensors and by the sensors already available in their environment.

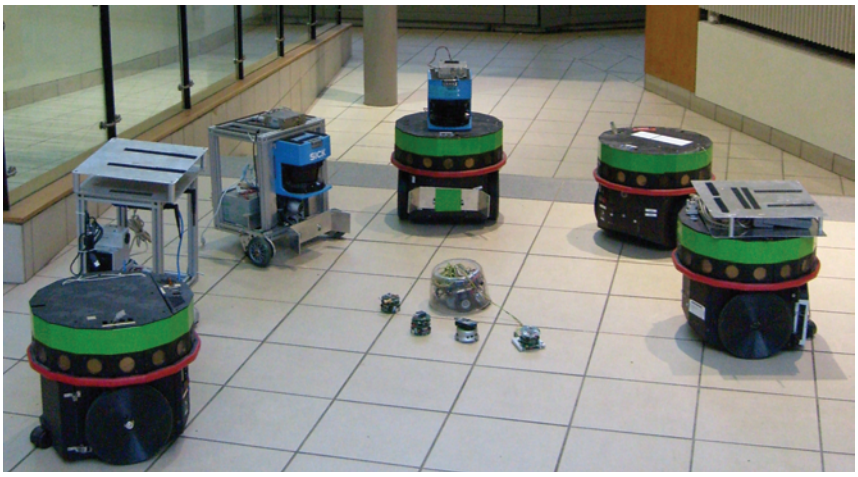

*Figure 1: Robots at the CLARITY Centre for Sensor Web Technology.*

Robots can also actively help the sensor network, not only by acting on the environment, but also by helping to deploy, program and maintain sensors; collaborating with sensors' localization; and by acting as mobile gateways in multi-hop networks. They also push the boundaries of WSN research by requiring interoperable and efficient solutions to data collection and online analysis.

## Adaptive and Self-Organising Software Architectures

The need for adaptive and self-organizing software architecture for ubiquitous robotics emerges from the very same requirements as for autonomous operations. In traditional networked robot systems, component integration is essentially an offline feature. In contrast, these systems demand a more open and dynamic approach, as the nature and availability of their hardware and software components are not stable but may change at run-time. In order to adapt to such environments, these applications must exhibit run-time flexibility, such as the ability to reorganize the interaction patterns of their architectural elements during execution. In particular, Component-Based Software Engineering (CBSE) and Agent-Oriented Software Engineering (AOSE) paradigms are natural candidates to provide modular architectures to integrate and dynamically organize the different system functionalities.

## Ubiquitous Personal and Social Assistant Agents

Not only will these robots have to deal with a variety of complicated tasks, but they will also be expected to behave in a socially intelligent and individualized manner in order to meet the diverse requirements of each user. However, reconciling the personalization/social aspect with pervasiveness and ubiquity remains a largely unexplored area of research. On both fronts, user interface agents, eg acting as a personal agent assistant (PAA) to their user, have already been widely adopted as intelli-

gent, adaptive social interfaces to the digital world, eg in the form of virtual characters interacting with the user via PCs, PDAs and the Internet. As such, the experience accumulated in these applicative domains may be used to inform robotics research. Moreover, as both software agents and robots increasingly inhabit the same human social space, their mutual interaction and combined operation within human societies will also acquire increasing importance.

### The New Testbed

The new testbed integrates and extends some pre-existing facilities, specifically:

- WSN of 70 Berkeley motes measuring humidity, light and temperature
- ten mobile robots, equipped with an array of state-of-the-art sensors, including USB cameras, laser range finders, sonar, infrared, odometers and bumpers. Each robot carries a mote able to measure ambient variables, which is also equipped with triple-axis accelerometers, magnetometer, compass and microphone
- a variable number of Internet gateways
- a variable number of PDAs and mobile phones equipped with Bluetooth.

Ubiquitous robot systems involve many interacting hardware and software components, and malfunctions in both types of component can encumber them. Such characteristics pose a considerable challenge to the application of a strong engineering perspective within robotics, as it is difficult to correctly gauge the accomplishment of specific system objectives and guide improvements and further developments. For these reasons, our research has commenced by defining a number of logging and inspection facilities.

On the robot side, an XML service developed in collaboration with the Interoperable Systems Group at DCU will provide a generic logging service for the instrumentation of both CBSE and AOSE frameworks. In addition to offline analysis, the service will also support hardware-in-the-loop simulations (HILS), thanks to synchronization and replay functionalities that will allow focusing online analyses and simulation over isolated groups of software modules.

The Octopus interactive dashboard, developed at UCD, will be employed to (i) visualize the topology and behaviour of the WSN through a network map; (ii) log collected parameters for network analysis; (iii) localize nodes on a provided 2D floor plan; and (iv) formulate and inject composite queries into the network.

**Links:**

CLARITY Centre for Sensor Web Technologies: http://www.clarity-centre.com/

CLARITY Ubiquitous Robotic Testbed: http://ubirobot.ucd.ie

IFR Statistical Department. 2007 world robotics survey: http://www.worldrobotics.org/index.php

The Interoperable Systems Group at DCU: http://www.computing.dcu.ie/~isg/

Octopus: http://www.csi.ucd.ie/content/octopus-dashboard-sensor-networks-visual-control

**Please contact:**
Gregory M. P. O'Hare
CLARITY Centre for Sensor Web Technologies, University College Dublin, Ireland
E-mail: gregory.ohare@ucd.ie

# Beyond RFID: The Ubiquitous Near-Field Distributed Memory

by Paul Couderc and Michel Banâtre

***The upcoming radio frequency indentification (RFID) revolution will undoubtedly contribute to the blending of the information society with the physical world: while common 'communicating objects' are currently restricted to complex electronic devices such as cell phones, cameras etc, RFID is in fact able to promote anything as a communicating object. In such a new world, data will not necessarily always flow into computers and networks, but may be physically retained by objects moving in real space.***

While RFID technology promises many useful applications, such as improved safety, easier and faster interactions, reduced error in data input and automation of tedious processes, it also raises serious concerns – in particular the privacy issue. In a world where many personal objects are electronically identified, the activities of individuals could be traceable in a similar, though much more comprehensive, way to 'googling' someone today on the Internet.

An important cause of this issue is that RFID systems are usually based on the concept of global identification, associated directory services or tracking databases. However, technically RFID are just small memory devices that can be addressed by near-field communication, and although identification has been the main application target, these devices provide support for alternative mechanisms.

We will now present such an alternative, through the example of the 'UbiCheck' service, a solution to a common problem when travelling: it is unfortunately quite easy to forget something. For example, security procedures in airports require that your personal effects are checked separately from you by X-rays. Forgetting one of your items, or mistakenly exchanging a similar item with someone else occurs frequently. Solutions have been proposed for this problem, based on active tags attached to the items that are monitored by an owner tag. This is impractical for several reasons: active tags are expensive, they require batteries (and hence regular maintenance), radio emissions may be restricted by regulations (on planes for example), and temporarily separat-

ing an item from its owner would require the alarm to be disabled.

Ubi-Check represents another solution using RFID tags attached to the items. It is possible to write in the memory of the tags the data required to check the integrity of the group of items. One possible implementation is to compute a digital certificate from the identifiers of all the items. An important aspect is that the identifiers associated with each item can be regenerated regularly (eg for each trip): they are only used for a locally computed integrity check, not for identification. The values could be written in the tags at, for example, the airport check-in, the train station, or even when leaving home. Then, at relevant points after the area in which people are vulnerable to item loss or exchange, we deploy checking gates (such as the exit of the security check in airports, or the exit gate of a plane or a train). These gates would ensure the integrity of groups of items crossing them, warning people in the case of a missing item or the presence of someone else's item.

This solution is a distributed system where only local properties are checked in order to ensure a global goal. In fact, it uses a principle similar to the transmission of a file in independent fragments over a packet network, where integrity is verified by checking sequence number coherency or checksums, except that here the data are carried by 'physical' fragments. Such a solution is interesting because while providing a security service, it avoids the privacy concerns of many other RFID approaches. Specifically, tracking of individuals is not easy, since the tags' content may change often for the same person and same set of objects. Further, the system is not based on identification, ensuring greater privacy.

Another interesting aspect is that checkpoints and association points are autonomous and only carry local processing. The system is therefore not dependent on a remote information system. This has important benefits in terms of extensibility, reliability and deployment costs.

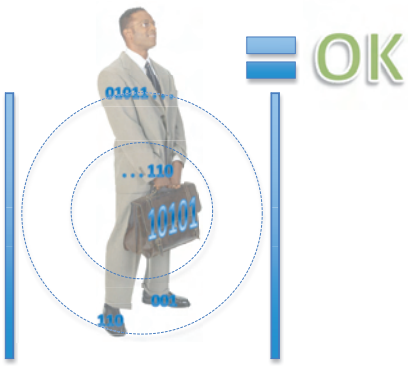


*Figure 1: Ubi-Check: integrity checking using a collection of eletronic tags*

This service is an example of a more general usage of near-field communication beyond identification. As mentioned previously, this technology consists of memory devices, a basic and generic support for computing, obviously not limited to storing identifiers. Another application example is to use RFID to build distributed maps for robotic support. In the Roboswarm project we consider very simple robots with limited knowledge and sensors. In order for them to navigate, RFID tags are arranged in the space where the robots are working. Each tag contains relative spatial references pointing to the nearest tags, effectively making a graph. These pointers allow a robot to reach one tag from another, assuming that odometry drift is low enough while moving the distance separating two tags. In this distributed map example, the memory is distributed quite thinly over the physical space. As near-field communication technology gets better and cheaper, denser and more ubiquitous memories are likely to emerge in the environment, with computing applications that we are yet to imagine. While the industrial promoters of RFID/NFC are mostly concentrating on applications based on identification, almost anything can be done with memory devices, including systems that respect privacy. A key question to making this possible is: will there exist completely 'blank' or free-format tags, without any pre-existing and non-rewritable identifier?

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

**Please contact:**
Paul Couderc
Centre de recherche INRIA Rennes - Bretagne Atlantique, France
Tel: +33 2 99 84 72 92
E-mail: Paul.Couderc@inria.fr

# Rapid Prototyping of Sensor-Based Applications with SunSPOTs

by Manfred Bortenschlager, Elisabeth Haid and Andreas Wagner

***Sensor technology has the potential to boost productivity just as the Internet did. We demonstrate the opportunities presented by sensors and sensor network technology by deploying a corresponding framework in an indoor environmental quality application. Our framework is based on OGC Sensor Web standards and exploits SunSPOT sensor technology for rapid prototyping.***

Due to the restrictions of industrial sensors, the engineering of applications that exploit sensor network technology is difficult. Such sensors are usually proprietary and inflexible as regards programmability, reusability and applicability to different domains. As a result, testing, extending and porting an application based on a sensor network is expensive. Java-based SunSPOT sensor technology has been developed for rapid prototyping of such applications and for testing and verifying algorithms on a small scale prior to deploying them in industrial operation. We demonstrate this methodology in the development of an indoor environment quality (IEQ) assessment application that exploits

SunSPOTs. This application is based on a wireless sensor network, adopts OGC Sensor Web standards, and provides a visual interface to the sensor values.

## SunSPOT Sensor Technology

SunSPOTs are embedded hardware modules that are equipped with a 180 Mhz CPU, 512 KB RAM, 4MB flash memory, three on-board sensors (temperature, light and three-axis accelerometer), hardware interfaces for the integration of arbitrary external sensors, and the IEEE 802.15.4 wireless transmission technology, which is enabled for mesh networking. SunSPOTs are entirely programmable in the Java programming language and thus help to abstract from the underlying hardware. No direct interaction by machine code with the hardware is necessary, which significantly eases the development of sensor-based applications.


*SunSPOT sensor.*

The SunSPOT sensor technology was developed primarily to satisfy three target groups: education (introducing pupils and students to related topics such as programming, networking, embedded systems, robotics or hardware), research and development (allowing for fast implementation of systems that integrate sensor technology and for easy testing of their behaviour), and hobbyists (developing small sensor-based applications for personal needs).

One of the major design goals of SunSPOTs was to provide a tool for rapidly prototyping sensor-based applications, and for testing and verifying algorithms on a small scale prior to deploying them in industrial operation. This tool can be exploited to reduce costs by evaluating novel algorithms a priori before adapting them to real-world problems.

This form of sensor technology will prove useful in a broad spectrum of applications and domains, such as early detection and warning systems, environmental monitoring, automotive engineering, warehouse/container management, logistics, monitoring of buildings, home automation, weather forecasting, medical monitoring of patients and diagnosis, and agriculture and farming.

## Application: Indoor Environmental Quality Measurement

Based on SunSPOT sensor technology we developed an indoor environmental quality application. Factors that influence the climate in a room and contribute to human well-being are measured at different positions in a room. The sensor data, then, can be retrieved from the sensors and are further processed in our sensor network.

Each sensor station is composed of two main components, namely (i) an external sensor which can measure electromagnetic pollution, air pressure, humidity, air temperature, brightness, noise or carbon dioxide, and (ii) a SunSPOT module which is responsible for preprocessing acquired sensor data and propagating them through the sensor network.

The base station managing this sensor network is an OGC-compliant Sensor Web application. It allows for administration of the sensors, and reading and processing of sensor data; for example, users can visualize and interaction with current sensor values on a graphical, Web-based interface. In particular, the Sensor Model Language (SensorML), Observation & Measurements (O&M), and Sensor Observations Service (SOS) specifications are adopted.

Communication between the sensor and the base station can occur in both push and pull modes and in a regular or on-demand fashion, where the values are communicated over the meshed wireless sensor network. All configurations can be defined by the user during runtime. In addition, a user can employ a sensor value reader device (essentially another SunSPOT) in order to get data from a specific sensor by physically moving into the communication range of that sensor and querying the accordant data.

## Future Activities

Further research will focus on the enrichment of the visualization component by correlating it with information potentially coming from other, external sources such as the Internet. This should be achievable with reasonable effort due to the adoption of open Sensor Web standards. In addition, methodologies will be investigated that will allow algorithms and functionalities that were successfully verified on a SunSPOT platform to be ported to non-Java sensor platforms with minimal effort.

The work presented here is part of the MobilityLab, a Centre of Excellence on the Engineering of Location-Based Systems, and involves Sun Microsystems, Vienna University of Technology (Institute for Geoinformation and Cartography) and Salzburg Research.

**Links:**

https://www.sunspotworld.com/
http://www.opengeospatial.org/projects/groups/sensorweb
http://mobilitylab.salzburgresearch.at

**Please contact:**

Manfred Bortenschlager
Salzburg Research Forschungsgesellschaft mbH / AARIT, Austria
E-mail: manfred.bortenschlager@salzburgresearch.at

# Building an Adaptive Environmental Monitoring System Using Sensor Web Technology

by Jer Hayes, Greg O'Hare, Harry Kolar and Dermot Diamond

***The ultimate goal of environmental sensor networks is to realize the concept of an 'adaptive environment' – one that senses and rapidly adapts to potential incidents in order to minimize their impact.***

Environmental pollution affects human health and reduces the quality of our land and water. As a result, there is great interest in monitoring water and air quality and ensuring that all areas are compliant with legislation. Ubiquitous environmental monitoring places considerable demands upon existing sensing technology. The combined challenges of system longevity, autonomous operation, robustness, large-scale sensor networks, operationally difficult deployments and unpredictable and lossy environments collectively represents a technological barrier that has yet to be overcome. The CLARITY Centre for Sensor Web Technologies is working with IBM to confront these challenges. Ubiquitous sensing envisages many aspects of our environment being routinely sensed. This will result in data streams from a large variety of heterogeneous sources, which will often vary in their volume and accuracy. The challenge is to develop a networked sensing infrastructure that can support the effective capture, filtering, aggregation and analysis of such data. This will ultimately enable us to dynamically monitor and track the quality of our environment at multiple locations. The ability to monitor quality is a prerequisite to maintaining quality, and ensures that detected pollution incidents are dealt with as quickly as possible, dramatically minimizing their impact. In effect, the ultimate goal of environmental sensor networks is to realize the concept of an 'adaptive environment' – one that senses and rapidly adapts to potential incidents to minimize their impact.

In recent years an increasing number of environmental incidents have occurred which would have benefited from an adaptive environment approach. These include the detection of lead in water supplies, cryptosporidium outbreaks and, as recently reported by the Irish EPA, and high pollution levels in a third of Irish rivers and streams. In Ireland and the rest of the world, proper management of environmental resources is critical for both health reasons and for the sustainable exploitation of these resources. Sensor Web technology such as low-power wireless communications, coupled with the emergence of new and reliable sensors such as the microfluidic analyser platform developed by CLARITY (see Figure 1), now enable environmental data to be collected at much higher temporal and spatial resolutions. Other examples of sensors being developed by CLARITY researchers are very low-power and low-cost colorimetric sensors that can be fabricated by modifying commercial optical components with chromo-responsive films. For example, inexpensive yet sensitive chemo-sensors can be made by applying such coatings to light emitting diodes (LEDs), as shown in Figure 2. Detailed experimentation with these devices in the field has demonstrated their effectiveness at detecting and tracking the dispersion of volatile chemical plumes. Continued improvement in the long-term reliability of chemical sensing platforms, coupled with significant reductions in their cost is the key to enabling scaled-up deployments of sensors at multiple locations, and this is a key goal for the CLARITY-IBM team.

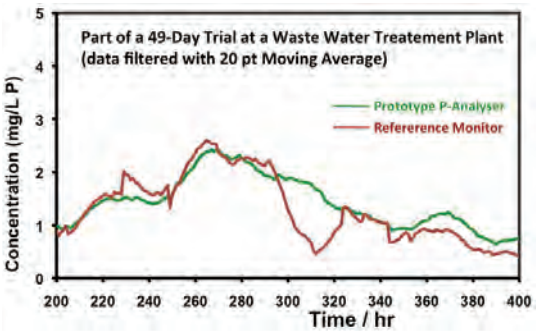


***Figure 1: The phosphate analyser developed by CLARITY (left) and results over an 8-day period from a field trial at a waste water treatment plant (right) showing excellent agreement with a reference analytical system (total duration 49 days).***

Sensor networks provide a web of interconnectivity: multiple sources of information that will allow decision-making processes to be more accurate and efficient. These processes can be complex and demanding however, and are often constrained in a number of possibly conflicting dimensions such as quality, responsiveness and cost. CLARITY researchers and IBM are working together to examine in-situ decision making, whereby decisions are effected based upon inferences made from both locally sensed data and data aggregated from sensor networks. The in-situ sensor nodes that comprise a sensor network are often computationally challenged with respect to processing power, as they are developed to be low-cost, low-power devices rather than high-performance computing devices.

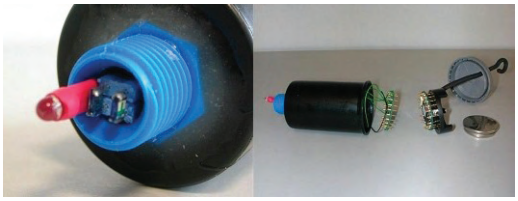

***Figure 2: A gas sensor made by coating a pH indicator dye (left) and the sensor platform, mica2dot mote (right). In the presence of an acidic vapour/gas, the sensor changes colour from deep blue to yellow.***

This computational challenge places a major constraint upon the reasoning process. To overcome this challenge, CLARITY and IBM are developing a hybrid reasoning approach to deliver in situ decision-making that combines stream-based computing with multi-agent system techniques. The hybrid reasoning approach builds on System S, an IBM technology for distributed stream processing, and Agent Factory, developed by CLARITY researchers to provide a powerful collaborative decision-making system. In partnership with the Irish Marine Institute, this hybrid approach is being tested through an environmental demonstrator project entitled SmartBay, which seeks to deliver an infrastructure to validate distributed in-situ real-time environmental monitoring technolgies.

SmartBay is a next-generation water management system (both marine and freshwater), aspects of which are currently under development by CLARITY for the Marine Institute of Ireland. It is located at Galway Bay on the Atlantic seaboard of Ireland. Collaboration in the development of SmartBay will form a core activity of the Global Centre of Excellence for Water Management, established by IBM in Ireland in 2008. SmartBay-related research will mirror the monitoring activities in the Hudson River in New York State of the Beacon Institute for Rivers and Estuaries, and will in part compliment the River and Estuary Observation Network (REON) established for that purpose.

The application of Sensor Web technology to environmental sensing will eventually result in the realisation of the 'adaptive environment' concept, though networks that can rapidly detect and adapt to potential environmental incidents, in order to minimize their impact. CLARITY and IBM are demonstrating key technological building blocks of the adaptive environment through projects such as SmartBay, through which new sensing and collaborative decision-making technology is being used to test elements of a networked sensing infrastructure that can support the capture, filtering, aggregation and analysis of environmental data. While the initial real-world benefits of such technology will be in the environmental domain, these benefits will ultimately belong to any application area of the Sensor Web.

**Links:**
CLARITY:
http://www.clarity-centre.com/
IBM Big Green Innovations:
http://www-03.ibm.com/technology/greeninnovations/
SmartBay:
http://www.marine.ie/home/services/operational/SmartBay/SmartBay.htm

**Please contact:**
Dermot Diamond
CLARITY Centre for Sensor Web Technologies, Dublin City University, Ireland
E-mail: dermot.diamond@dcu.ie

Harry Kolar
IBM, T.J. Watson Research Center, USA
E-mail: kolar@us.ibm.com

# Remote Water Monitoring With Sensor Networking Technology

by Thiemo Voigt, Nicolas Tsiftes and Zhitao He

***Sensor networks enable remote monitoring of natural environments such as glaciers, volcanoes and bodies of water. Within the project 'Sensor Networks to Monitor Marine Environment with Particular Focus on Climate Changes', SICS and partners are designing and implementing flexible, reprogrammable sensor network solutions suitable for monitoring the marine environment with high resolution in time and space.***

Carrying out marine research requires that studies be undertaken in remote environments such as the Baltic Sea. However, marine environmental monitoring is expensive: the cost of operation at sea includes at least €10 000-15 000 per day for the use of a lar ge research vessel, plus the cost of laboratory technicians, analytical instrumentation and logistics. As a result, the marine environment is poorly monitored. For example, in the Baltic Sea's Bothnian Bay, the nine fixed stations are visited only eight times a year.

Sensor networks mean that marine data collection can be undertaken in a much more cost-efficient fashion. Within the multi-disciplinary project 'Sensor Networks to Monitor Marine Environment with Particular Focus on Climate Changes', SICS and partners have designed an advanced water monitoring system (Figure 1). Our system features a diving unit, consisting of a bin containing several sensors connected to one sensor node. The diving unit moves up and down an anchor line. In this way we are able to take measurements at any depth using only one set of sensors. A pressure sensor determines the depth at which measurements are being recorded. When the diving unit is not in the water taking measurements, it is parked in a garage that is part of our buoy construction. By parking the diving unit in the garage, we expect to eliminate fouling, ie the accumulation of material on hard surfaces in aquatic environments. We first intended to use an oil-filled garage, but found that this affected the sensor readings. Initial experiments with an air-filled garage have shown that this solution also prevents fouling.

Our system is driven by rechargeable batteries. Previous projects have had problems with solar power. Our project partner Uppsala University therefore designed a wave energy generator to obtain electric power from vertical wave movements. As a ring of stacked magnets floats up and down along a rod wound with metal coils, an alternating voltage is induced across the two ends of the coil. This is further rectified and

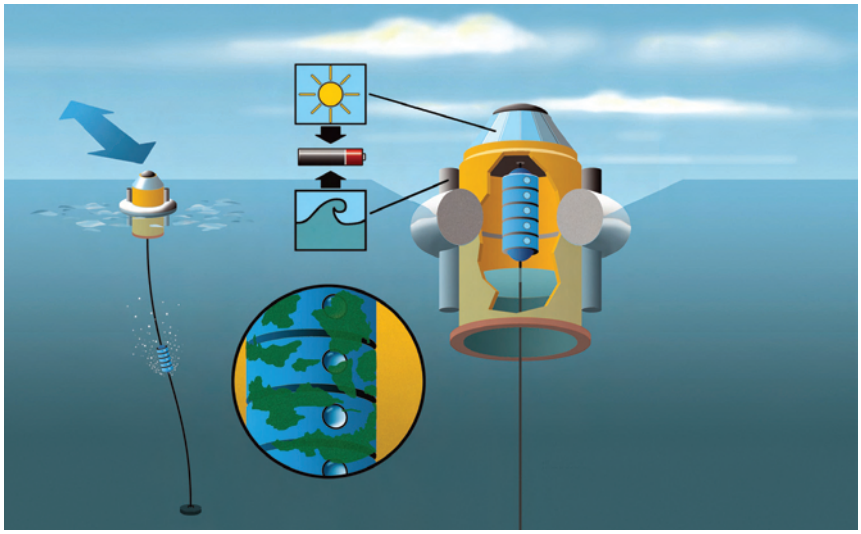

Illustration: Bo Rönnerdahl

*Figure 1: Sensor system for remote water monitoring.*

filtered to provide a constant charging voltage to the batteries.

Our system is designed to cope with water depths of around 100m, meaning the number of measurements taken during one dive can be significant. Storage is thus required for the large volumes of measurement data: not only those collected by the diving unit but also from light and acceleration sensors placed on the top of the buoy. System data also needs to be saved; for example, the energy produced by the wave energy generator as well as the available energy. The sheer amount of data exceeds the limited RAM available in sensor nodes. Moreover, to handle the different data types smoothly, we use Contiki's Coffee file system to store data. The data is transferred onshore using GPRS (General Packet Radio Service), and this is too energy-consuming a task to be performed after every dive. We have therefore ported Coffee to operate on SD cards that enable cheap mass storage.

Due to the high energy consumption of the GPRS unit as well as the per-byte cost of the transferred data, we compress the data before sending it. Previous work has shown that compression is much cheaper than communication for typical sensor nodes with low-power radios. The same is true for the more power-hungry GPRS communication. To this end, we have designed a new compression algorithm called SB-ZIP that is more efficient than state-of-the-art compression algorithms. For example, SB-ZIP compresses 4.5 MB of acceleration data measured on the buoy to only 1.5 MB, whereas the well-known S-LZW algorithm is only able to reduce it to slightly over 3 MB.

SICS is running this project in collaboration with Umeå Marine Sciences Centre, Uppsala University and SMHI, the Swedish Meteorological and Hydrological Institute. At the conclusion of the project in late 2009, we expect to have deployed around ten buoys at places relevant for marine monitoring in the Baltic Sea.

**Link:**
http://www.sics.se/node/1361

**Please contact:**
Thiemo Voigt, SICS, Sweden
Tel: +46 8 633 1598
E-mail: thiemo@sics.se

# Sensors Anywhere – Sensor Web Enablement in Risk Management Applications

by Gerald Schimak and Denis Havlik

***The increasing frequency, severity and consequences in Europe of floods, storms, forest fires and other natural hazards sensitive to climate change has clearly shown the shortcomings of existing environmental monitoring and information systems. The observed inefficiency is primarily a consequence of historical and organizational factors. An exorbitant amount of work on data and service standardization would be required to build more efficient information systems using state-of-the-art technology.***

Emerging technology in risk monitoring and management has the potential to speed up the necessary organizational and structural changes. Ad hoc wireless sensor networks and the collective intelligence of the Sensor Web; the plug-and-measure paradigm of IEEE 1451 smart sensors; the Semantic Web; and the OGC Sensor Web Enablement architecture: all these address critical factors of the state-of-the-art technology. However, none of the currently available and emerging technologies offers rapid deployment, easy maintenance, quality assurance and automated data processing along the whole information processing chain from smart sensors and wireless ad hoc sensor networks, over automated data loggers and value-added middleware services, to user applications capable of dynamically integrating all available data sources at run time.

Sensors Anywhere (SANY) is an ambitious FP6 IST Integrated Project dealing with sensor networks for environmental and risk management applications (Figure 2). SANY aims to contribute to joint efforts of the European Commission (EC) and the European Space Agency (ESA) on 'Global Monitoring for Environment and Security' (GMES) by improving the interoperability of in situ sensors and sensor net-

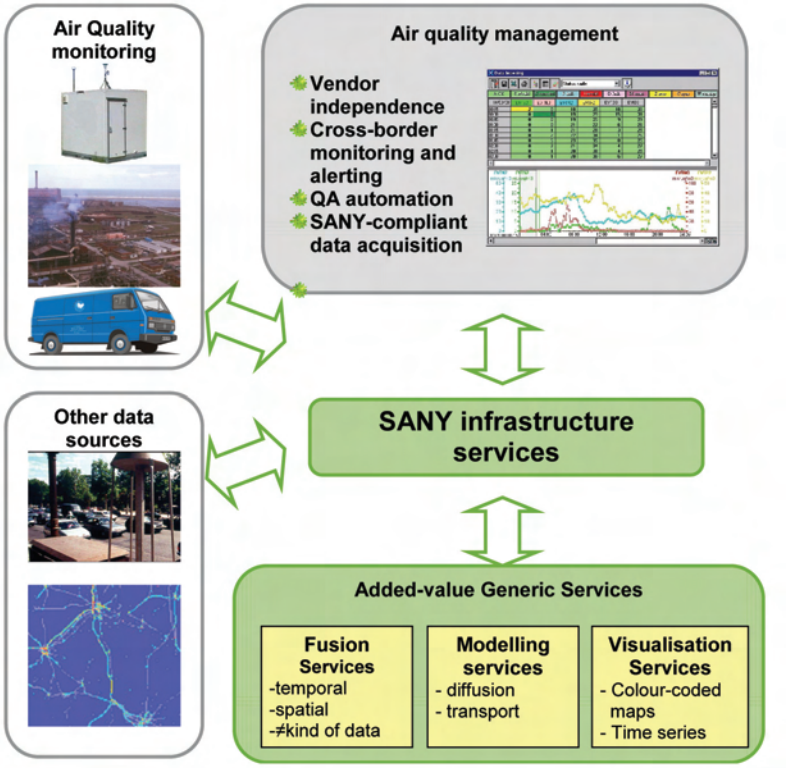


*Figure 1: SANY air quality management subproject.*

works and taking up the challenges mentioned above. This means specifically that the Sensor Service Architecture (SensorSA) developed in SANY shall allow quick and cost-efficient reuse of data and services from currently incompatible sources in future environmental risk management applications across organizational, administrative or regional borders.

In order to assure the sustainability of the project results, SANY reuses the open standards from W3C, OASIS, ISO and Open Geospatial consortium (OGC). One of the most promising standardization efforts is currently happening within the Open Geospatial Consortium Sensor Web Enablement initiative (OGC-SWE). The goal of OGC-SWE is to enable all types of Web and/or Internet-compatible sensors, instruments, and imaging devices to be accessible and, where applicable, controllable via the Web. The OGC-SWE vision is to define and approve the standards foundation for interoperable Web-based sensor networks. For this purpose, OGC-SWE specifications include a standardized model for representing and exchanging observation results as well as an information model and encodings that enable the discovery and tasking of Web-resident sensors.

Of special interest for SANY are the following: a service for retrieving sensor 'observation' data and meta-information, the so-called 'Sensor Observation Service' (SOS); a service for sensor planning and executing tasks, called the 'Sensor Planning Service' (SPS); a service that allows users to subscribe to specific alert types, known as the 'Sensor Alert Service' (SAS); and a service that facilitates asynchronous message interchange between users and services, and between two OGC-SWE services, called the 'Web Notification Service' (WNS).

Three validation subprojects of SANY demonstrate the feasibility of building risk management applications with SensorSA.

*Figure 2: 'SANY'-trees - sensors anywhere in our environment.*

The first validation subproject (Figure 1) illustrates the use of SensorSA services at data acquisition level, as well as vendor independence and the feasibility of building virtual networks across administrative and technical borders. Furthermore, it is concentrating on assessing air pollution episodes, tracking pollution back to its source and predicting the air pollution in urban areas or around industrial zones.

The second validation subproject concentrates on coastal water management issues, such as assessing, modelling and predicting bathing water quality with user-determined 'what-if' scenarios. It uses advanced data fusion services to demonstrate the feasibility of run-time binding of heterogeneous data sources and the run-time choice of data fusion service, and to predict microbiological contamination through data fusion using all available real-time measurements and the results of microbiological water sample analyses. The real-time data sources used by this subproject include meteorological data, sea temperature, turbidity, salinity, chlorophyll and dissolved oxygen levels.

Finally, the third validation subproject concentrates on geohazard monitoring in a complex urban environment, which means assessing the structural instability of architectural objects caused by human activities (eg tunnel excavation). It demonstrates the integration of wireless ad hoc sensors into SensorSA networks; rapid deployment, auto-configuration and self-management; and data fusion of in situ and Earth Observation data sources.

The authors wish to thank the Information Society and Media Directorate General of the European Commission (DG-INFSO) for co-funding the integrated project SANY within the area of ICT for Environmental Risk Management.

**Link:**
http://www.sany-ip.eu

**Please contacts:**
Denis Havlik
Project co-ordinator S@NY
Austrian Research Centers – ARC (AARIT)
Tel: +43 50550 3157
E-mail: denis.havlik@arcs.ac.at

# FLASH: Fine-Grained Localization in Wireless Sensor Networks using Acoustic Sound and High-Precision Clock Synchronization

by Evangelos Mangas and Angelos Bilas

***Sensor localization is an important component and enabler of many applications using intelligent sensors. FLASH is a system that achieves fine-grained localization using acoustic sounds and high-precision clock synchronization via radio frequency (RF) communication. Sensors can dynamically localize themselves in space by maintaining synchronized clocks and measuring time of arrival for acoustic sound pulses. Our techniques focus on achieving highly accurate synchronization and consistent sound detection. Furthermore, FLASH does not require external infrastructure such as fixed equipment, specialized hardware support, or great resource consumption. Experimental results show that FLASH localization is accurate to within 11cm in a variety of indoor environments.***

Localization is an important component of many sensor applications. For instance, the ability to accurately locate humans or objects in enclosed spaces can help in detecting and responding to abnormal situations. Techniques such as the use of GPS, image processing or static ranging mechanisms are either not effective, cumbersome to use, or incur high costs.

Our work focuses on achieving fine-grained localization without excessive or specialized resources that use wireless sensor nodes. We rule out the use of extensive external infrastructure, since it is usually both expensive to acquire and complex to install. Moreover, it is important to achieve localization by using only typical resources available in sensor nodes, without the use of specialized hardware peripherals. For these reasons, we choose an approach that relies on an acoustic sound-ranging scheme. Each reference node produces an audible sound pulse while the rest of the nodes are sensing the audio frequency spectrum. All listening nodes attach a timestamp to the received sound pulse, in a global synchronized timescale. These (synchronized clock) timestamps are used to calculate the time of flight for each sound pulse. Distance is then estimated based on the speed of sound. Our work is therefore focused on high-accuracy clock synchronization and sound detection for range estimation.

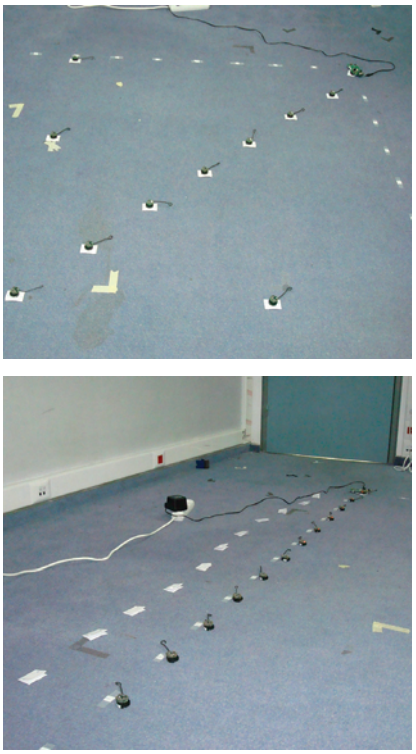

*Figure 1: Performing experiments with sound in our lab.*

To achieve precise clock synchronization we implement a synchronization protocol operating at the Media Access Control (MAC)-Layer that does not introduce significant communication costs and that deals with fixed overheads introduced by the interrupt mechanism and RF communication. We also introduce an external mechanism for testing synchronization precision, causing simultaneous interruptions to synchronizing motes and then comparing the synchronized timestamps produced. We implement our synchronization protocol on Mica2dot motes. Our experi-

***Figure 2:***
***Left: Error in estimated distance with respect to absolute distance;***
***Right: Localization in two dimensions, the sounders are node0 and node1.***

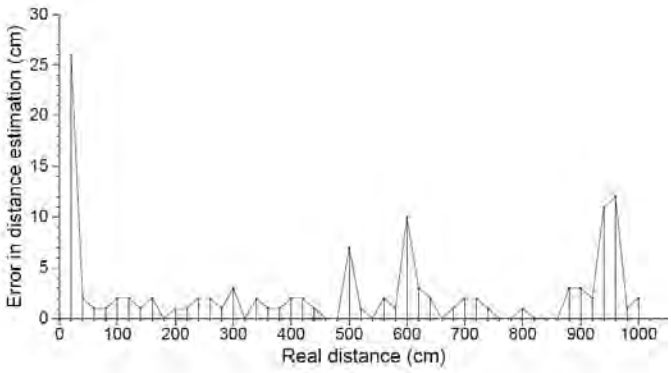


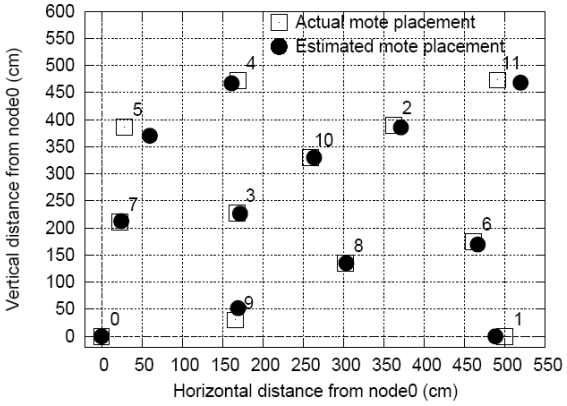

mental results show that synchronization precision has an error of less than 5µs (median) for a 30s resynchronization period.

Our acoustic sound technique focuses on consistently detecting the start of the arriving sound pulse. Attaching a timestamp at the beeper node is simple: right before applying voltage to the embedded buzzer we use the local (synchronized) clock to broadcast an RF message. At each listening node we identify a sound pulse produced by the beeper node by using the periodicity of the sound pulse and requiring the average peak-to-peak amplitude to surpass a certain predefined dynamic threshold, which is less susceptible to reflections. Then we attach a timestamp to the first peak-to-peak measurement that was of greater value than our dynamic threshold. The difference between the two timestamps at the beeper node and the listening node results in the time of flight for the sound pulse.

Overall, FLASH demonstrates that localization can be achieved with inexpensive off-the-shelf devices and yet be quite precise. Figure 1 shows part of our experimental setup, whereas Figure 2 shows our results for 1D and 2D localization. Using the Berkeley Mica2dot motes, the default microphone that already exists on the sensor platform and cheap simple buzzers, we were able to locate nodes at distances of up to 10m, depending on surrounding noise. The average error in localization precision is 11cm for distances up to 7m. However, our approach does not require either calibration or any special infrastructure. Furthermore, our method requires a single sounder and microphone per node, resulting in better energy efficiency compared to methods that use multiple sensor devices per node.

Future work in the area includes localization in outdoor and more demanding environments, such as in the presence of obstacles between nodes, intense noise in a room, temperature and humidity variations and outdoor environments. We believe that FLASH and similar techniques will play an important role in cyber-physical systems and in our efforts to better interact with and control our environment.

**Link:**
http://www.ics.forth.gr/carv/scalable/

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

# High-Density Wireless Geophone Networks for Oil and Gas Monitoring and Exploration

by Stefano Savazzi, Vittorio Rampa and Umberto Spagnolini

***Strong fluctuations in crude oil prices and the expected production peak of current reservoirs are pushing oil companies to increase their investment in seismic exploration. Replacing cabling with wireless technology should radically improve the quality of depth imaging and simplify acquisition logistics. Recent advances in Wireless Sensor Networks (WSN) now allow the wireless community to satisfy the rigid constraints imposed by seismic acquisition systems, which have a large number of sensors (> 10 000) over the monitoring area (> 5km$^2$).***

Strong fluctuations in crude oil prices are pushing oil companies invest more in seismic exploration of new oil reservoirs and in new technology to improve the quality of depth imaging. Seismic prospecting requires a large number of sensors (up to 30 000), such as geophones or MEMS-based (Micro Electro-Mechanical Systems) accelerometers. These are deployed over large areas (up to 30km$^2$) to measure the back-scattered wavefield generated by an active excitation source. A storage/processing unit (sink node) collects measurements from all the geophones in real time to obtain an image of the sub-surface. Current telemetry is cable-based and usually requires hundreds of kilometers of cabling, which results in delays, high logistic costs and low imaging quality.

Wireless technology is thus expected to significantly improve the efficiency of oil exploration. Technical limitations in the data-rate efficiency, interference and battery use of current short-range wireless network architectures (eg WiFi, Bluetooth) forced previous proposals for wireless geophone system architectures to choose a combination of wireless and wired configuration. However, recent advances in WSN technology conveniently address the issues related to the strong constraints imposed by seismic acquisition systems. A Wireless Geophone Network (WGN) must support multiple acquisition settings and applications. Basic network requirements are: i) network throughput of 150kbps down to 50kbps for single component sensors; ii) real-time (or near real-time) acquisitions with strong delay constraints; iii) remote control by sink node and synchronous acquisition with a maximum timing skew of 10µs; and iv) accurate positioning of each sensor/geophone with an error of less than 1m to avoid degradation of the depth imaging quality.

## Network Architecture

As shown in Figure 1, the proposed WGN architecture exploits different radio transmission technologies to efficiently handle both short-range transmissions (ie for short-distance low-power communication among geophones/sensors), and long-range transmissions (ie for seismic data delivery to storage units and geophone remote monitoring) that must cover distances of several kilometers. The hierarchical network design requires the deployment of a number of Wireless Geophone Gateways (WGGs) to collect data readings from a large number of wireless geophones (WGs) and forward the data to the storage unit (SU). These WG nodes are self-organized into independ-

ent sub-networks; ideally the number of devices per sub-network should be as high as 300 nodes to minimize the number of WGGs. This results in an aggregated (per sub-network) throughput of about 45Mbps (up to 60Mbps). Data delivery within one sub-network is obtained by multi-hop transmissions towards the WGGs; WG sensors are within 5-100m of inter-node distance to reduce both energy consumption and increase battery life.

### Physical and MAC Layer Requirements

The requirements of self-localization and frame synchronization make Ultra WideBand (UWB) technology the natural choice for short-range transmissions within each sub-network. To achieve positional accuracy with errors less than 1m, the travel-time estimation error for ToA-based (Time of Arrival) positioning must be in the order of 3ns with a minimum required signal bandwidth of 500MHz. UWB technology provides data acquisition, synchronization and localization without the use of fully GPS-based (Global Positioning System) WGN nodes. Moreover, recent advances in radio design (ie MB-OFDM or MultiBand Orthogonal Frequency Division Multiplexing) provide wireless devices with high data rates over short ranges of up to 480Mbps, and low power consumption (ie below 100mW in active transmission mode but down to 20μW in power-save mode). The MB-OFDM processing can also guarantee network scalability through time and frequency division by allowing the use of multiple sub-bands to separate the co-located sub-networks, and coexistence with other 2.4 GHz-based radio devices without significant cross-interference.

The high number of devices per sub-network and the large network size suggest the adoption of a number of distributed MAC (Medium Access Control) functionalities. Network topology should define a hierarchical structure where the WGG acts as an intermediate sink towards the storage unit. The WiMedia standard (ECMA-368 from ECMA International, the European Association for Standardizing Information and Communication Systems) has been chosen as the reference for the development of the WGN MAC. Transmission is organized in superframes with the beacon period (BP) carrying the essential information of each device. Logical device/sensor groups are dynamically formed according to WiMedia protocol to facilitate the sharing of resources, while wireless medium reuse can be exploited over different spatial regions.

WGG supports specific extended functions compared to a standard WiMedia device. These functions allow: i) the Gateway to behave as an intermediate sink, forwarding data to the storage/processing node SU and controlling each sub-network; ii) contention-free resource negotiations to guarantee real-time constraints (eg quality of service and maximum delay); and iii) coexistence of long/short range transmissions. Figure 2 illustrates the MAC layer framing structure adopted for each sub-network, while the probability of full network coverage versus the BP length is shown at the bottom of the same figure. Sensors/geophones are assumed to be deployed according to the requirements of a conventional seismic survey. Geophone deployment has a major impact

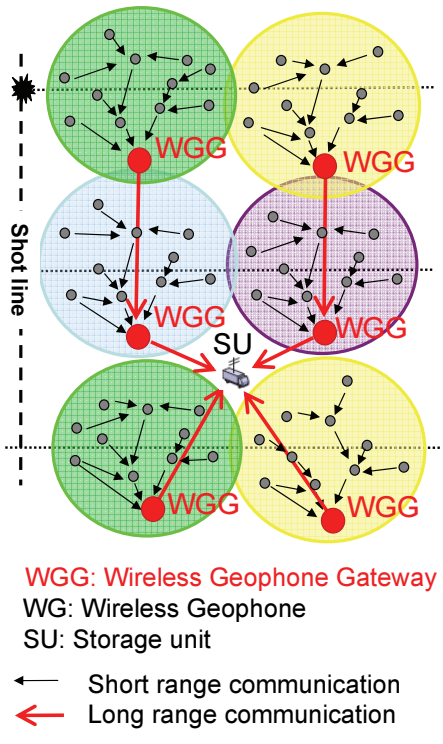


*Figure 1: Wireless Geophone Network architecture.*

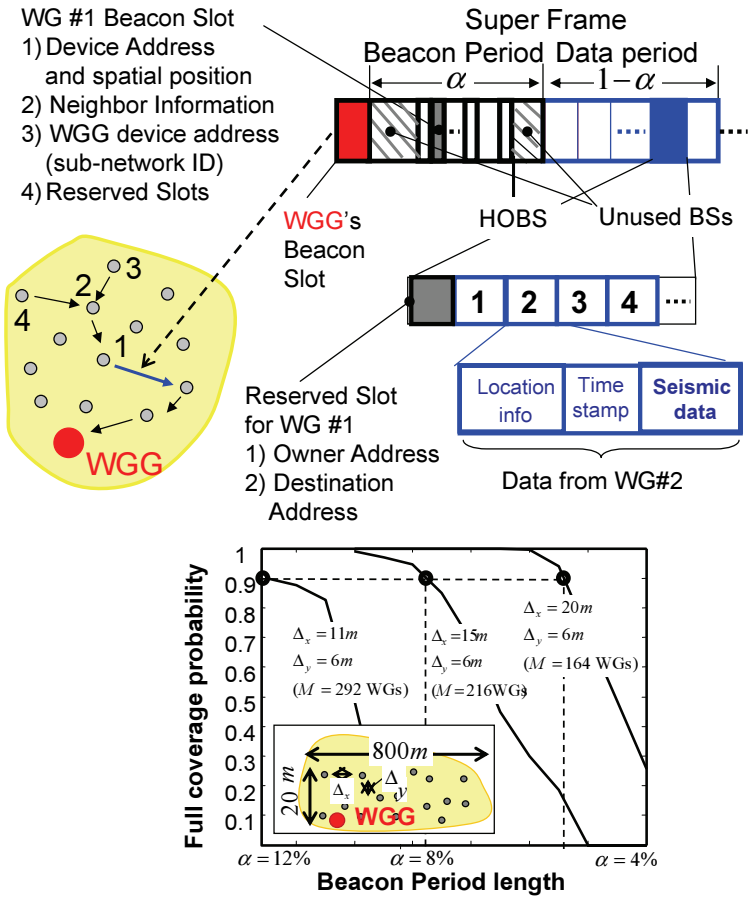


*Figure 2: MAC layer framing structure (top) used by each sub-network, and its impact on full network coverage (bottom).*

on the framing structure design: a lower density with a sensor spacing of 20m (still reasonable for seismic acquisitions) can make WiMedia feasible for WGN applications with minimal modifications. On the other hand, higher geophone densities, with a sensor spacing of less than 11m, require further MAC modifications. More details are given at the WisyGeo Web site.

**Links:**
Land seismic exploration:
http://www.oilfieldreview.com/

Wireless Geophone Network:
http://www.wisygeo.com
http://scitation.aip.org/tle

WiMedia:
http://www.wimedia.org

**Please contact:**
Stefano Savazzi, Umberto Spagnolini,
DEI, Politecnico di Milano, Italy
E-mail: savazzi@elet.polimi.it,
uspagnolini@elet.polimi.it

Vittorio Rampa
IEIIT-CNR, Italy
E-mail: rampa@elet.polimi.it

# A Software Platform for the Acquisition and Online Processing of Images in a Camera Network

by Thomas Sarmis, Xenophon Zabulis and Antonis A. Argyros

***Applications related to vision-based monitoring of spaces and to the visual understanding of human behaviour, require the synchronous imaging of a scene from multiple views. We present the design and implementation of a software platform that enables synchronous acquisition of images from a camera network and supports their distribution across computers. Seamless and online delivery of acquired data to multiple distributed processes facilitates the development of parallel applications. As a case study, we describe the use of the platform in a vision system targeted at unobtrusive human-computer interaction.***

Camera networks are increasingly employed in a wide range of Computer Vision applications, from modelling and interpretation of individual human behaviour to the surveillance of wide areas. In most cases, the evidence gathered by individual cameras is fused together, making the synchronization of acquired images a crucial task. Cameras are typically hosted on multiple computers in order to accommodate the large number of acquired images and provide the computational resources required for their processing. In the application layer, vision processing is thus supported by multiple processing nodes (CPUs, GPUs or DSPs). The proposed platform is able to handle the considerable technical complexity involved in the synchronous acquisition of images and the allocation of processes to nodes. Figure 1 illustrates an overview of the proposed and implemented architecture.

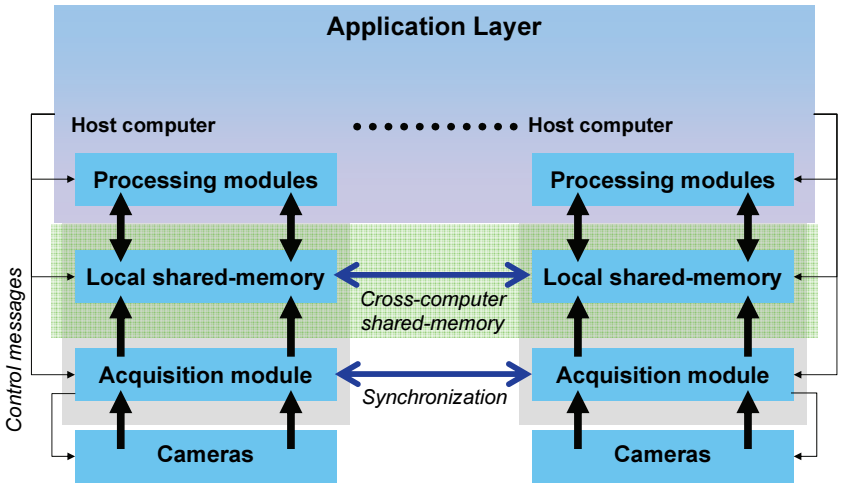


*Figure 1: Platform architecture for camera network supported by multiple computers. Central sensor control is provided by a message-passing communication infrastructure. Acquisition modules place the images on a synchronized shared-memory space, making them available to multiple processing nodes across computers. Processing modules read these data and perform parallel execution of computer vision algorithms. The same modules also have access to this space, allowing them to synchronously view intermediate computation results.*

The platform integrates the hardware and device-dependent components employed in synchronous multi-camera image and video acquisition. Pertinent functionalities become available to the applications programmer through conventional library calls. These include online control of sensor-configuration parameters, online delivery of synchronized data to multiple distributed processing nodes, and support for the integration and scheduling of third-party vision algorithms.

System modules can communicate in two modes. Communication through message-passing addresses control messages to targeted or multicast recipients. The diversity of communicated information types is accommodated by data-structure serialization. Communication through shared-memory spaces provides visual data or intermediate computation results to the nodes of the host or of multiple computers. The large bandwidth requirements imposed by image transmission are accommodated by a Direct Memory Access channel to a local shared-memory space. For cross-computer availability of images, memory spaces are unified over a network link. The latency introduced by this link is compensated for by notification of nodes, regarding the partial or total availability of a synchronized image set. In this way, per-frame syn-

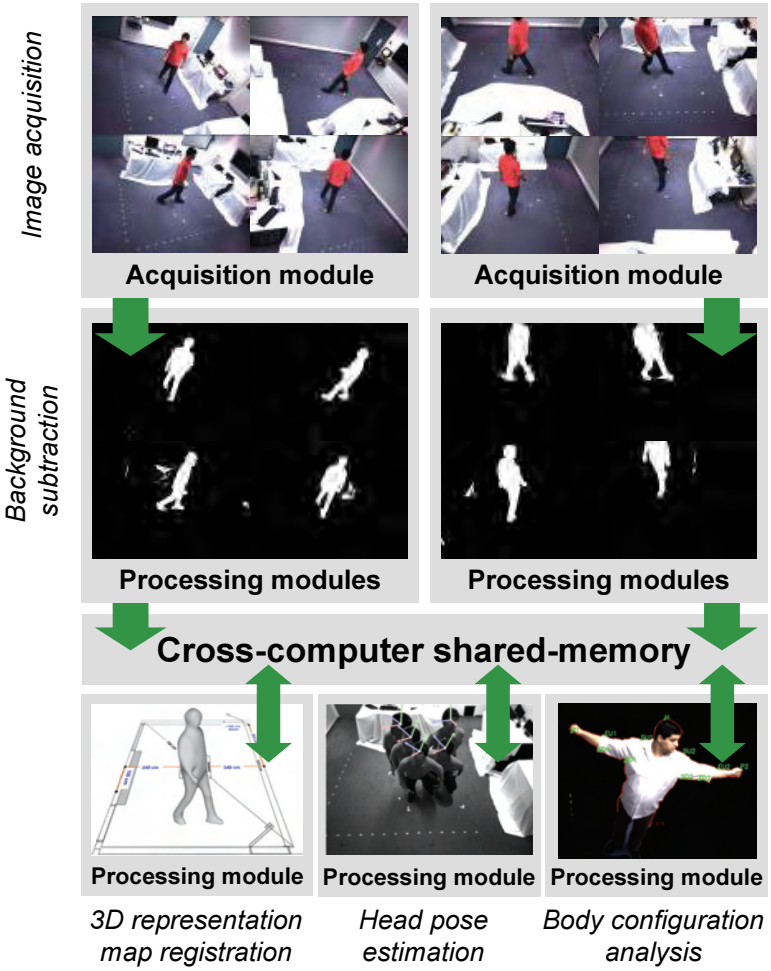


*Figure 2: Person silhouettes in synchronous images are segmented in parallel. Individual processes fuse them into a registered 3D representation of the person, recognize coarse gestures as expressed by body configuration, and estimate the spatial direction in which the person's head is facing.*

chronization of modules is achieved, but at the same time, processing of partially available input is also supported. Shared-memory spaces across processing nodes are essential, as large data capacity and frequent input rate demand the parallelization and pipelining of operations.

Acquisition modules encapsulate the complexity of sensor-specific, synchronization, and shared-memory configurations. Online sensor configuration and command is implemented through message-passing, while image transmission utilizes shared-memory communication. A range of off-the-shelf sensor types is supported through an extensible repository of device-specific wrappers. To facilitate testing of applications, input may be prerecorded.

Processing modules run vision algorithms that are transparent to the computer and provide access to images and intermediate computation results. During the applications development stage, an Application Programming Interface enables synchronization and message coordination. Articulated application development is facilitated by support for 'chaining' of processes.

Being in the format of a binary library, this platform can be invoked, independent of the programming language used. As an additional utility, the developed platform provides a GUI for the control of generic camera networks and the recording of image sequences. Forthcoming extensions involve additional capabilities for cooperation with middleware infrastructures in systems where vision is integrated with other sensory modalities (aural, tactile etc).

The platform is currently employed in the development of a vision system (illustrated in Figure 2), targeting unobtrusive and natural user interaction. The development of this system is part of a broader project funded internally at FORTH-ICS on Ambient Intelligence (AmI) environments. The system employs multiple cameras that jointly image a wide room. Two computers host eight cameras and a dedicated bus for their cross-computer synchronization, and utilize a LAN connection for communication. Upon image acquisition, a sequence of image processing and analysis operations is performed in parallel on each image, to detect the presence of humans through background subtraction in the acquired images. Using the shared memory across computers, segmentation results are fused into a 3D volumetric representation of the person and registered to a map of the room. Two other processes run in parallel and access the same data to recognize the configuration of the person's body and estimate the pose of the person's head. The utilization of the proposed platform facilitates the modular development of such applications, improves the reusability of algorithms and components and reduces substantially the required development time.

This work has been partially supported by the FORTH-ICS RTD programme 'AmI: Ambient Intelligence Environments'.

**Link:**
http://www.ics.forth.gr/cvrl/miap/doku.php?id=intro

**Please contact:**
Xenophon Zabulis
FORTH-ICS, Greece
E-mail: zabulis@ics.forth.gr

# Tackling the Semantic Gap in Multimodal Sensor Networks

by Eric Pauwels, Albert Salah and Paul de Zeeuw

***Sensor networks are increasingly finding their way into our living environments, where they perform a variety of tasks like surveillance, safety or resource monitoring. Progress in standardization and communication protocols has made it possible to communicate and exchange data in an ad hoc fashion, thus creating extended and heterogeneous multimodal sensor networks. CWI is looking at ways to automatically propagate semantic information across sensor modalities.***

Wireless sensors are deployed in a growing number of applications where they perform a wide variety of tasks. Although this has considerable economic and social advantages, it seems likely that even greater benefits can be gained once heterogeneous sets of individual sensors are able to communicate and link up into larger multimodal sensor (inter)networks. We expect that the network's performance will become more robust when information from multiple sources is integrated. Indeed, an additional layer of intelligence on top of the communication protocols will enable sensors to advertise their own capabilities, discover complementary services available on the network and orchestrate them into more powerful applications that meet high-level specifications set by human supervisors. This can be achieved more efficiently if the capabilities of the different components can be described in both human- and machine-readable form. It will then be possible for individual sensors to relate their own objectives and capabilities to human-defined goals (eg minimize energy consumption without sacrificing comfort) or available knowledge, both of which are usually expressed in terms of high-level semantics.

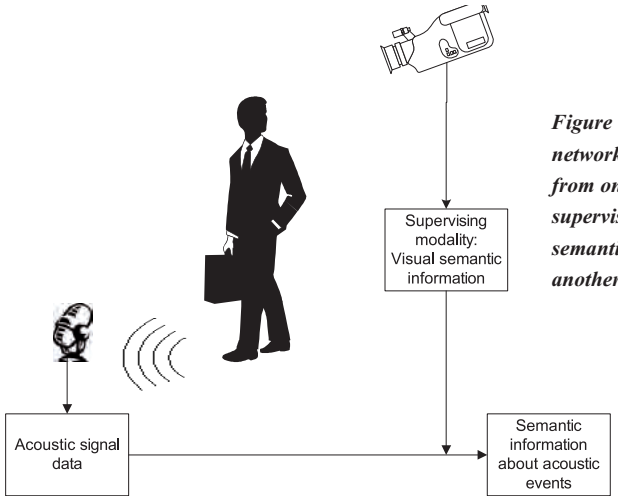


***Figure 1: In multimodal sensor networks, reliable information from one sensor can be used to supervise the extraction of semantic information from another sensor.***

In addition, networks could become smarter for at least two reasons: sensors that produce highly reliable output can be used to provide on-the-fly 'ground truth' for the training of other sensors within the network, and correlations among sensed events could bootstrap the automatic propagation of semantic information across sensors or modalities.

The implementation of our vision requires two conditions to be met. Firstly, sensors should come equipped with an open interface through which their output data and all relevant metadata can be made available for third party applications. Secondly, sensor networks need to be endowed with a learning mechanism that shifts the burden of supervision from humans to machines.

Granted, the linking of low-level sensor data to high-level semantic concepts remains a formidable problem, but we contend that the complementarity inherent in the different sensing spectra supported by such a network might actually alleviate the problem. The basic idea is simple: if particular sensor data can be linked to specific semantic notions, then it can be hypothesized that strongly correlated data picked up by complementary sensors (or modalities) are linked to semantically related concepts. A simple example will clarify the issues at hand: imagine a camera network on a factory floor that has been programmed to identify persons using face recognition, and to determine whether or not they are walking, say for safety reasons. If the same factory is also equipped with open microphones that monitor ambient noise, then an intelligent supervision system might pick up the strong correlation between walking people as observed by the camera network and rhythmic background sounds as detected by the microphones.

By mining general knowledge databases, the system might then be able to conclude that the observed rhythmic audio output corresponds to the sound of footsteps and add this snippet of semantic information to its knowledge database. In essence, the system succeeded in using available high-level information (the visual recognition of walking people) to bridge the semantic gap for an unrelated sensor (audio). By accumulating the information gleaned from such incremental advances, we contend that it will be possible to gradually - but largely automatically - extend the system's knowledge database linking low-level observed sensor data to high-level semantic notions.

To explore the viability of this idea we have conducted a number of simple experiments in which we used the Internet as a general knowledge database. For instance, referring to the above scenario we submitted the paired search terms walking (as the camera has been programmed to detect this behaviour) and sound (through the use of standards such as SensorML, each sensor can communicate the modality of its output) into a search engine and analysed the response. By restricting attention to

meaningful words that occur frequently (both in terms of number per page and number of unique pages), we end up with a sorted list that suggests a link between the audio data and a list of semantic concepts including music, video, gait, work and footsteps. In a final step this list is further whittled down by checking each of these suggestions against an ontology to determine their semantic distance to the original concept (walking). By restricting attention to the most similar concepts, it transpires that it is highly likely the recorder audio is related to either footsteps, gait or music, all of which make sense. These results hint at the possibility of automatically extending semantic notions across modalities, thus leading to more robust and intelligent networks.

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

**Please contact:**
Eric Pauwels
CWI, The Netherlands
Tel:+31 20 592 4225
E-mail: Eric.Pauwels@cwi.nl

# Utilising Wearable Sensor Technology to Provide Effective Memory Cues

by Aiden R. Doherty and Alan F. Smeaton

***We describe a wearable sensor technology that passively records 'lifelog' images and sensor readings of a wearer's daily life. The focus of our work is not on aggregating, collecting or networking data as in the usual application of sensors in the Sensor Web, but rather on detecting events of interest to the wearer from a multi-sensor standalone device. These events of interest provide effective cues to allow people to more easily access their autobiographical memories. Early research indicates this technology may be potentially helpful for sufferers of neurodegenerative diseases such as Alzheimer's.***

Sensors and sensing technology are everywhere, and this issue of ERCIM News contains many examples of sensors networked together for some greater purpose. Mostly, people deploy sensors and then gather the readings together and address issues like networking, calibration, sensor fusion and sensor event detection. The general trend is towards networking sensors into the Sensor Web, but this isn't the only way of using them. Sensors can be used in small groupings on standalone devices that gather and process information and report back not sensor readings, but major semantic events. In this article we describe one such sensor technology which is simple and cheap to manufacture, but can empower an individual to reflect on their past behaviour and memories.

Lifelogging is the term used to describe the recording of different aspects of your daily life, in digital form, for your own exclusive personal use. It can take many forms, such as an application running on your mobile phone that 'logs' all your phone calls. One particularly interesting device is the SenseCam, a camera that is worn around the neck and automatically captures thousands of images of the wearer's life every day. It has a range of in-built sensors for monitoring the wearer's environment, detecting movement, ambient temperature, passive infrared information (ie body heat) and light intensity.

Preliminary studies indicate that information gathered by the SenseCam is potentially useful as a memory aid to recall autobiographical memories. Research in the field of cognitive neuropsychology has established that 'cued recall' is better than 'free recall'. The closer a cue is to how an actual memory was encoded, the better memory retrieval is. Other studies indicate that autobiographical memories tend to be strongly encoded in a visual manner in the brain. The SenseCam records pictures from the viewpoint of the user, making it able to provide visual cues of our past that are very close to how the original memories/experiences were encoded in the brain.

*Figure 1: Sample lifelog images from a typical IT worker's day.*

Even though SenseCam images provide strong memory cues, there exists a substantial problem in effectively managing the overwhelming volume of images generated by this device – approximately 650 000 images per year are captured. Within the CLARITY centre at Dublin City University, we have developed a suite of functions applied to SenseCam data that automatically provide effective digital memory retrieval cues. We structure our processing into a number of logical steps that exploit various characteristics of the human memory system.

1. Firstly, we intelligently segment sequences of images into distinct events such as having breakfast,

working on a computer etc. This is achieved very quickly using on-board environmental sensor values.

2. Given that human memory stores information associatively, we provide users with automated search functions to find events similar to a given event, eg "show me other times when I was at the park". By intelligently representing events through the fusion of image descriptions and the in-built sensor values, we found that users can find events related to any given 'query event'.

3. Given that the human memory more strongly encodes distinctive memories, we automatically identify events that are more visually unique among those recorded by a wearer. We have found that it is effective to combine the automated detection of faces (to indicate social engagement) with detecting how visually novel each event is.

4. As human memory is known to store items associatively, it is useful to augment individuals' SenseCam events with images (or videos) from external sources, eg to better remember a trip to the Eiffel Tower by viewing pictures of the tower uploaded by others to the Internet. Using GPS information, and after some intelligent automated processing, we can automatically find relevant supplementary images and videos from Internet sites such as Flickr and YouTube.

Our technology for autobiographical memory capture and management has been deployed not only within our research centre in Dublin, but also in numerous cognitive psychology research groups in Europe and North America, and uses sensors and a camera to gather information as part of a lifelog. These include Universities of Toronto, Tampere, Illinois, Utrecht and CWI in Amsterdam. Our approach does not conform to the common model of a sensor network composed of inter-connected sensors with live, real-time streaming data. This is because the demands of lifelogging are for post-event reflective retrieval rather than real time, meaning live inter-connectivity with other sensor nodes is not as vital as in other Sensor Web technologies.

**Link:**
http://www.cdvp.dcu.ie/SenseCam/

**Please contact:**
Aiden R. Doherty, Alan F. Smeaton
CLARITY Centre for Sensor Web Technologies, Dublin City University, Ireland
E-mail: adoherty@computing.dcu.ie, alan.smeaton@dcu.ie

# SENSE – Smart Embedded Network of Sensing Entities

by Wolfgang Herzner

***The SENSE project (Smart Embedded Network of Sensing Entities) is developing methods, tools and a test platform for the design, implementation and operation of smart adaptive wireless networks of stationary embedded sensing components. The network is an ambient intelligent system, which adapts to its environment and delivers reliable information to its component sensors and the user.***

SENSE is an EC-funded project of the 6th Framework Programme, Embedded Systems (objective 2.5.3, contract no. 33279). It aims at developing a platform for smart adaptive wireless networks of smart sensors. These sensors cooperate to establish and maintain a coherent global view from local information. Newly added nodes automatically calibrate themselves to the environment, and share knowledge with their neighbours. The network is scalable due to the local processing and sharing of information, and self-organizes based on the physical placement of nodes.

As test platform for a civil security monitoring system, a test application composed of video cameras and microphones, was chosen. The test platform will be installed at the Krakow Balice airport, to yield real data and performance goals from a realistic test environment. Each sensor is a stand-alone system consisting of multiple embedded components: video system, audio system, central processor, power source and wireless networking. The security application implements object/scenario recognition (eg unattended luggage or people 'lurking' in an area). Nodes recognize local objects, using a combination of video and audio information, and neighbouring nodes exchange information about objects in a self-organizing network. The result is a global overview of current objects and events observed by the network (see Figure 1).

The five main objectives are:
- to build networked systems of embedded components that can dynamically and automatically reconfigure themselves
- to convert low-level local information to semantic knowledge
- to use semantic-level knowledge for network-centric computation
- to understand how a shared semantic vocabulary influences dynamic node discovery and configuration
- to understand how perception and information processing can be combined using low-level and high-level feature fusion.

The expected results of SENSE are to combine the aspects of:
- embedded intelligent middleware in smart devices,
- adaptive configuration,
- flexible cooperation (among devices),
- high-level perception and adaptation
- dynamic networking in a common framework of semantic knowledge discovery and sharing.

The SENSE system encompasses aspects including:
- construction of a modality-neutral embedded test platform
- raw sensory processing

- transformation of sensory data into semantic knowledge
- communication between nodes to produce a consistent world view
- sharing of knowledge between intelligent nodes
- automatic recognition of unusual and alarm situations
- communication between the intelligent network and an operator, and
- automatic discovery and configuration of new intelligent nodes.

Embedded systems in SENSE develop their own semantic symbols based on an analysis of their environment. SENSE incorporates research from machine learning to discover statistical regularities in its environment, and compresses these regularities into informative semantic symbols. At the local level, SENSE uses algorithms such as 'expectation-maximization' to optimize each node's set of semantic symbols. Sharing of knowledge between nodes is also a topic of research, both in distributed systems and artificial intelligence. The SENSE system uses a mature algorithm called 'belief propagation'. This algorithm specifies how to share probability distributions over semantic concepts between nodes, such that a self-consistent world view results. Figure 2 illustrates the architecture of SENSE.

The unique feature of SENSE is that it combines technology from embedded systems, robotics, networking and machine learning research in a new way. The result is a framework for the development of smart networks of embedded components that are flexible, adaptive and device-independent. Networks that cover those challenges are called ad hoc networks or self-organizing networks. Their development is driven by the wireless community, but some of their principles are also of interest for wired networks. Less effort is required for their installation, initialization and maintenance, and they display inherent fault tolerance and the possibility to save energy within the network. This is true for both wired and wireless types but is typically only relevant for the latter. In contrast to self-organizing networks, traditional networks have a very time-demanding commissioning phase that also involves expert knowledge.

Dynamic addition or removal of nodes is a further challenge. In ad hoc networks, connections are constantly created and destroyed. This is called 'plug and participate'.

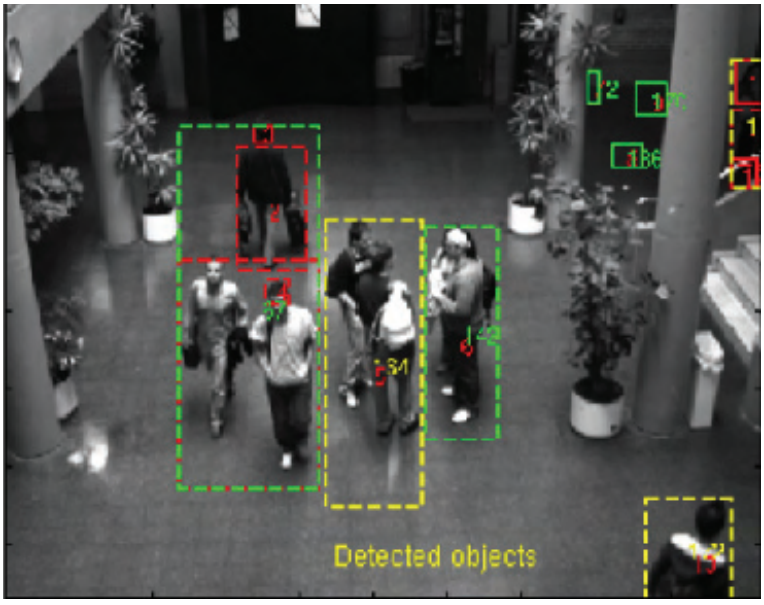


*Figure 1: Civil Security demonstrator (Airport Krakow); frame-by-frame detection of objects.*

Vision and audio sensors were selected to fit the application domain and also because they are complex enough to allow for significant advancement in sensor processing technology. However, the framework designed will be generic enough to accommodate a wide range of sensors. The middleware and software framework is designed to easily incorporate additional sensor types, which mean the project results are more easily extended and the technology is easier to adopt for third parties.

With respect to requirements, the airport application is very similar to other applications in public buildings and public areas, such as shopping centres, railway and bus stations, office buildings, football stadiums and so on. The SENSE technology is modular and is appropriate for integration with existing security systems (existing video systems, for example). This provides an upgrade path from current systems to a full distributed SENSE system, improving the chances for technology uptake.

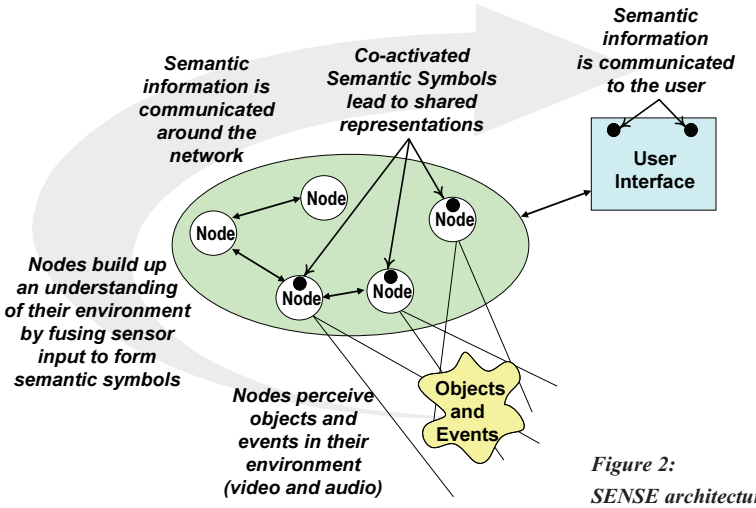


*Figure 2: SENSE architecture.*

**Link:**
http://www.sense-ist.org

**Please contact:**
Wolfgang Herzner
Austrian Research Centers – ARC (AARIT)
E-mail: wolfgang.herzner@arcs.ac.at

# Towards Data Management in the Sensor Web: the MaD-WiSe System

by Giuseppe Amato, Stefano Chessa, Francesco Furfari, Stefano Lenzi, Claudio Vairo

***The convergence of sensor networks with the Web (Sensor Web) poses new problems. These relate both to the management of the enormous amounts of data continuously produced by sensors, and to the reaction to events inferred from such data. The MaD-WiSe system (MAnagement of Data in WIreless SEnsor networks), developed at ISTI-CNR, exploits the well-known database paradigm to address this issue.***

The MaD-WiSe system considers a Wireless Sensor Network (WSN) as a highly distributed and dynamic database. Sensor data can be acquired, manipulated and filtered using simple SQL-like statements. Various WSN applications can be developed using its functionality. For example, the MaD-WiSe system has been used to implement a prototype application providing remote monitoring of firefighters equipped with totally encapsulated chemical suits (see Figure 1). This application enables real-time monitoring of various physiological parameters during operational activity. The information acquired can be used to raise the alarm in situations of risk, provide advice to the team leader with respect to the necessary actions, and also to analyse offline the health of those involved in an operation. The application employs five sensors that are deployed on the arms, legs and chest of the operator.


*Figure 1: Firefighters being monitored with MaD-WiSe.*

The MaD-WiSe architecture comprises a set of modules running on WSN nodes (network side), and a set of modules that offer WSN services to the external applications (context information provider). The network side consists of a set of modules that implement a distributed data stream management system on a WSN. It is organized into three layers, as shown in Figure 2. The layers interact through well-defined interfaces and are autonomous with respect to each other.

The Network Layer supports connection-oriented multi-hop communication between arbitrary pairs of nodes. The Stream System Layer offers abstraction mechanisms for data access by means of data streams. It can be thought of as the equivalent of a file system on a sensor network. The main difference is that in the Stream System, data is continuously acquired from transducers, communications between nodes and data processing.

The Stream System defines three types of streams: sensor, remote and local streams. A sensor stream is connected to a transducer and carries data originating from the transducer. A remote stream is a data channel between two distinct sensors: writing to a remote stream occurs on one sensor while reading from the stream occurs on the other. Thus remote communication between different sensors is encapsulated within the stream system, which in this respect offers the transport layer functionalities. A local stream is local to a sensor, as writing to and reading from the stream can only be requested by code running on the same sensor.
The Stream System allows streams to be created or removed, and records to be read from and written to existing streams. Data rates can be associated with both sensor and remote streams. For the former, data rates determine the activation frequency of transducers associated with sensor streams. In the latter case, data rates are used by the network layer to optimize radio scheduling: the radio is switched on only when a piece of data must be sent through a remote stream. Sensor streams can also be 'on demand'. In this case the transducers are activated only in response to an explicit read request on the stream.

The last of the three layers is the Query Processor Layer, which implements the query processor of a full-in-network distributed data stream management system. It can be programmed by the client-side subsystem in order to take part in the execution of a distributed query. The query language used in MaD-WiSe is called MW-SQL and shares its basic constructs with SQL. However, sensor

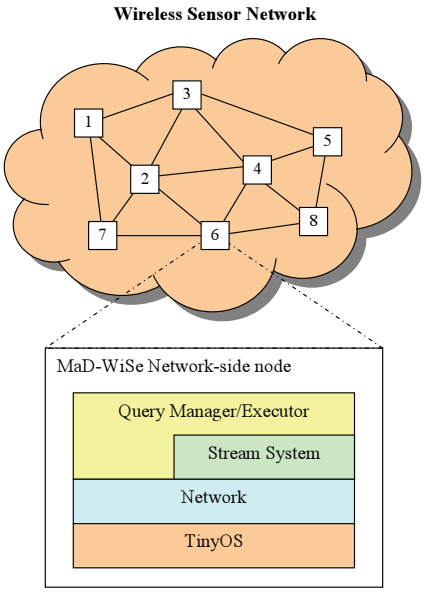


*Figure 2: The software in a network-side node.*

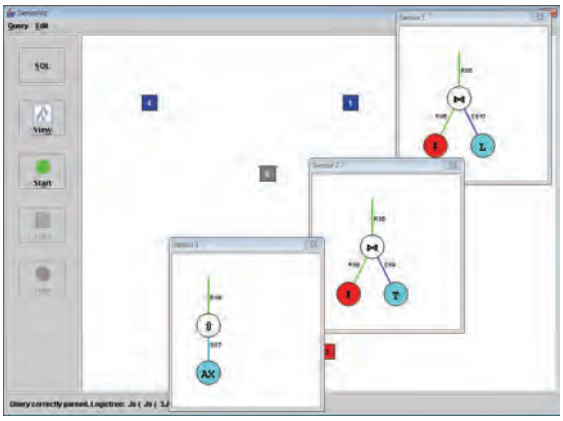

*Figure 3: The MadWise query interface.*

network peculiarities and the distributive nature of the database implementation introduce some differences. MW-SQL allows users to express queries to manipulate, filter and organize sequences of tuples generated by the sensors. MW-SQL relies on the concept of source to present the user with an abstraction of a sequence of tuples arriving from a precise origin. The MaD-WiSe query interface is shown in Figure 3.

The MaD-WiSe context information provider fits within the raw data retrieval layer of a context-aware architecture. It comprises a low-level module (composed of a query parser, an execution plan optimizer and a query manager) and a higher-level module, the JDBC driver, which interacts with the low-level module by means of the MW-SQL language. At the current stage of the project, the JDBC driver is being encapsulated within an OSGi bundle in order to implement a gateway between the WSN and the Sensor Web, and to enable queries involving different WSNs in the Sensor Web.

MaD-WiSe was developed using nesC and is available for both TinyOs 1.x and 2.x. It runs on WSN platforms based on MICAz and IRIS motes and is distributed with an open-source licence. Further information and downloads can be found on the project Web site.

**Link:**
http://mad-wise.isti.cnr.it

**Please contact:**
Giuseppe Amato
ISTI-CNR, Italy
Tel: +39 050 3152906
E-mail: giuseppe.amato@isti.cnr.it

# Detecting Hazardous Gases in Emergency Disaster Scenarios using Wearable Sensors

by Tanja Radu, Cormac Fay, King Tong Lau and Dermot Diamond

***The aim of this project is the development of integrated smart wearable sensors for emergency disaster intervention personnel. The CLARITY (The Centre for Sensor Web Technologies at Dublin City University) team is involved in the integration of gas sensors into wearables for detection of hazardous gases like CO and $CO_2$.***

This research arose through involvement in a joint European Union-funded FP6 project called Proetex (www.proetex.org). The aim of the project is to develop textile- and fibre-based integrated wearable sensor systems. Such systems will improve the safety and efficiency of emergency personnel by monitoring the health status of the operator and the surrounding environment for potential risk sources. Some of the issues covered are monitoring of vital signs, posture and activity, external hazard monitoring, and low-power wireless communications. The project is closely connected with other large European projects on smart textiles, wearable sensing and associated applications – it brings together and extends the technology developed by previous projects (see for example www.biotex-eu.com).

The project brings together 23 partners from a range of backgrounds – universities, research institutions, industrial partners and end users – drawn from laboratories in France, United Kingdom, Poland, Italy, Belgium, Switzerland, Germany and Ireland.

Researchers from CLARITY based in the National Centre for Sensor Research at Dublin City University are involved in the integration of sensing platforms into wearables for the detection of environmentally harmful gases surrounding emergency personnel. Special attention is being paid to carbon monoxide (CO) and carbon dioxide ($CO_2$). These gases are associated with fires and mining operations, and it is of the highest importance to warn and protect operators from potential harm caused by over-exposure to high concentrations of these gases. The objective is rapid detection of the status of an environment (low, medium or high hazard) and real-time communication of this information to the garment wearer. Critical in this identification of potential toxification is a reliable method of measuring $CO/CO_2$ exposure. Commercially available sensors have been carefully selected and are being integrated into the outer garments of firefighters. The sensors provide sufficient sensitivity to reliably alert users to the presence of these harmful gases. Another important aim is to achieve wireless transmission of sensor signals to a wearable wireless base station that gathers, processes and further transmits the data.

When selecting the appropriate commercially available sensors for the gas sensing application, special attention was paid to sensor size, robustness, sensitivity and power requirement. Electrochemical sensors satisfy most of these requirements, especially in terms of size and power requirements. CO is detected using an amperometric sensor

*Figure 1: Firefighter's boot with built-in pocket used for enclosing the $CO_2$ sensor and wireless communications platform.*

in which the current between the electrodes is proportional to the concentration of the gas. On the other hand, the $CO_2$ sensor is potentiometric. In this case, the reference and working electrodes are placed in an electrolyte that provides a reference $CO_2$ concentration. The measured potential is based on the difference in concentration between the reference electrode and the outside air. Both types of sensors are very sensitive and give an accurate reading (in parts per million). This means that both low concentrations of these gases (which can be hazardous over long periods of exposure) and high concentrations (which pose an immediate danger) can be accurately detected. The signal obtained from these sensors is transmitted wirelessly to the wearable base station using Zigbee. Power is supplied to the sensors using a nickel metal hydride rechargeable battery. The $CO_2$ sensor is placed in a specially designed pocket located on the firefighter's boot. The pocket is designed not to obstruct the firefighter's activities. The prototype currently used for testing is shown in Figure 1; note the side pocket containing the $CO_2$ sensor along with the wireless sensing module and a battery. The pocket has a waterproof membrane that protects the sensor from moisture, but allows gas to pass through. The CO sensor will be integrated in the firefighter's outer garment (ie jacket). All sensed information will be fed to a wearable local base station that shares the data with a remote centralized base station. The ultimate goal is to achieve local communication between firefighters and civil workers in the operations area, as well as longer-range communications between these personnel and the support team outside the operations area.

The project commenced in 2006 and will end in early 2010. The project is envisioned to produce three sets of prototypes during the four years of its existence; so far two generations of prototypes have been successfully developed. At this stage, accurate wireless transmission of the sensor signal has been successfully achieved (see Figure 2). Future activities will include evaluation of prototypes in laboratory and field conditions. Their performance will be compared to that of existing technology, and will be improved upon by customizing the products according to the specific user needs. Finally, the products will be tested in real-life situations as an ultimate proof of their full functionality.

The authors gratefully acknowledge the financial support of the European Union (Proetex FP6-2004-IST-4) and the Science Foundation Ireland (07/CE/I1147).

**Links:**
http://www.proetex.org
http://www.dcu.ie/chemistry/asg/radut
http://www.clarity-centre.com

**Please contact:**
Tanja Radu
CLARITY Centre for Sensor Web Technologies, Dublin City University
Tel: +353 1 700 7602
E-mail: Tanja.Radu@dcu.ie

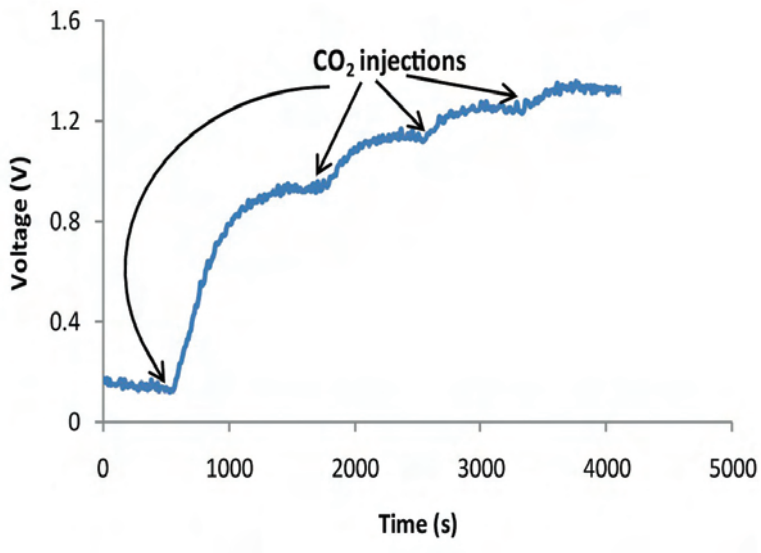


*Figure 2: Wirelessly transmitted signal from $CO_2$ sensor calibration (range atmospheric to 42000 ppm CO2). Sensor was enclosed in an airtight chamber and $CO_2$ was injected.*

# TennisSense: A Multi-Modal Sensing Platform for Sport

by Noel E. O'Connor, Philip Kelly, Ciarán Ó'Conaire, Damien Connaghan, Alan F. Smeaton, Brian Caulfield, Dermot Diamond and Niall Moynahan

***The evolution of the World-Wide Web to the Sensor Web is providing an unprecedented opportunity to develop novel applications in a variety of domains. In this article we describe our work on multi-modal sensing for sport.***

Sensor technology is rapidly changing the professional sporting landscape. Modern motor racing has been virtually transformed by the introduction of sophisticated drive-by-wire sensor technology, for example. Similarly, the ability to accurately monitor the performance of an athlete during training is having a major influence on a wide range of track and field events. In this project we introduce state-of-the art sensing technology onto the tennis court with a view to facilitating coaches as they train the next generation of tennis superstars. These sensors make it possible for coaches to obtain a second-by-second record of player performance that goes far beyond what can be captured by more traditional techniques; simply eyeballing the player as they move, serve, volley and return, does not reveal what might be going on under the surface in terms of the player's ever-changing performance profile, his/her breathing patterns, and the stresses and strains that they are feeling in their joints and muscles.

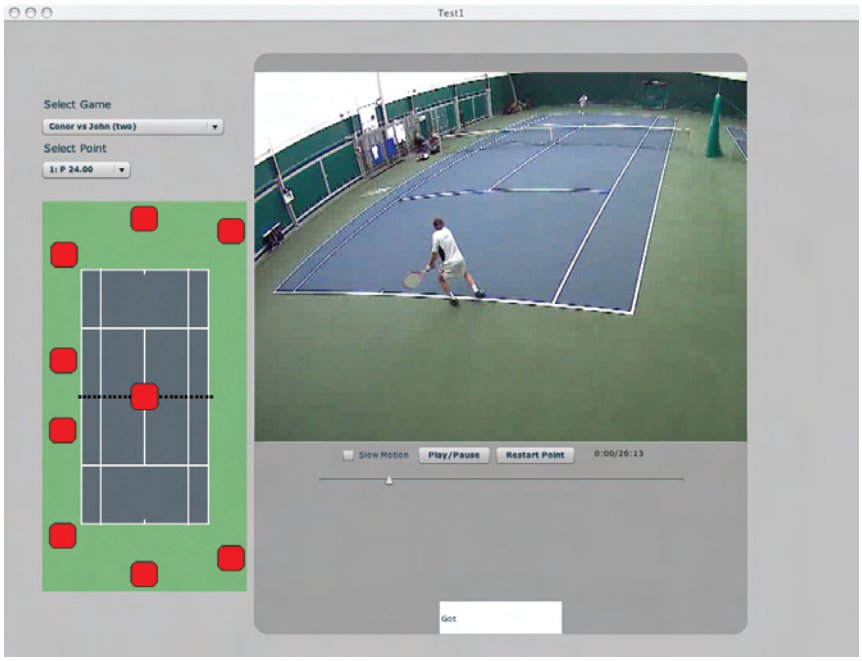


*Figure 1: Coaching tool developed for Tennis Ireland.*

The project is a collaboration between CLARITY (The Centre for Sensor Web technologies) and Tennis Ireland, the national governing body for tennis, based in Dublin City University. We have instrumented an all-weather tennis court with nine Internet-enabled cameras with built-in microphones. This is linked to a localization system that identifies the player's position to within 15cm by triangulating the radio signal emitted by small tags carried by the players in their pockets. The coach uses a simple wireless device to signal an important event during play via a simple button press. After training, these button presses are synchronized to the video streams. Content analysis mechanisms are then used to define the precise start and end of tennis play around these locations. We do this by tracking the ball in each camera, measuring player speed using the location information and detecting the sound of a ball hitting a racquet. These detected events are used in a variety of Web-based coaching tools. These include (a) being able to view individual tennis strokes from multiple view points, (b) support for online coaching feedback (eg "This shot was the wrong choice given the location of your opponent when you played it"), and (c) an easy way for coaches and athletes to produce personalized annotated video summaries for players to download and review for motivational purposes.

## Integrating Smart Materials and Body Sensor Networks

By its very nature, this project requires convergence between multiple disciplines in order to be successful. It requires input from engineers for signal processing and content analysis aspects and from computer scientists for addressing indexing, archival, personalization and user interface issues. Sports scientists play the crucial role of mediating between the technologists and the end-user coaches and athletes. Their input is invaluable in helping translate a complicated set of domain-specific and expert-driven requirements into a set of concrete technical functionalities. This helps ensure the practical relevance and usefulness of any technology developed. They also help ensure that coaches and athletes understand the potential benefits of the technology, thereby stimulating take-up by the key stakeholders.

However, the potential for convergence extends even beyond this initial constituency. More generally, the sensed environment can be considered as an experimental platform for trialling wearable sensing. The wireless inertial sensing platforms being developed by CLARITY engineers will be integrated into the system with a view to providing biomechanical feedback to athletes.

The smart materials being developed by chemists in the CLARITY work programme will be integrated into body sensor networks and will allow a variety of important physiological and biometric indicators to be detected.

To date, most data gathering and experimentation with this technology happens within a laboratory. The Tennis Ireland installation provides infrastructure for data gathering, analysis and feedback that will allow this technology to be tested in the field, during real exercise. For example, the work on conducting polymers will result in textiles for measuring stretching, bending and pressure movements. This 'smart foam' will be built into the wearable platform of one of our industry partners and trialled with tennis players to monitor specific muscle movement and strain. Data collection and feedback can be piggy-backed on the existing system, with the added benefit of a synchronized multi-view video to augment expert analysis of data from the wearable sensors.

A longer-term example is the 'sweat patch' being developed to measure the pH of sweat. A change in sweat pH may signal an increased reliance on anaerobic metabolism, but there is a lack of reliable methods for assessing pH during exercise. CLARITY is working on a new fluid-handling platform based on polyamide lycra. The device consists of a super-absorbent material that provides a passive pumping mechanism to control fluid flow. The optical detection of pH-induced colour changes in the dye is achieved via a paired emitter-detector LED system. The Tennis Ireland installation provides an ideal framework for testing this innovative technology in a real application scenario with real users.

**Please contact:**
Noel E O'Connor
CLARITY Centre for Sensor Web Technologies, Dublin City University, Ireland
Tel: +353 1 700 5078
E-mail: oconnorn@eeng.dcu.ie

# Synchronizing Sensed Data in Team Sports

by Dónall McCann, Mark Roantree, Niall Moyna and Michael Whelan

***In this article we will be discussing the synchronization of sensor data in team sports. Synchronization allows us to use more expressive queries, to query across all participants in a given activity and to potentially discover new knowledge from the semantically enriched data. A collaborative research effort between groups working on data management and on health and human performance (both at Dublin City University) involved a series of experiments using wearable sensors during team games and the capture and querying of sensed data.***

When dealing with sensor data for a team sport, it is often useful to be able to query across multiple sensors and thus to be able to compare data from several players for any given moment in time. In order to do this, the data from all sensors must be synchronized so that the start time of the game or activity can be identified in the data from each individual sensor.

This is necessary because sensor devices may be activated asynchronously, since the device begins recording when it first comes into contact with the player's skin. While many sensor devices will record a start time, this information is not necessarily reliable as there is often no correlation between the system time and the time kept by the match officials, or indeed between the times on any two sensors. In addition, the devices may be unreliable and may malfunction, or the device may become detached during the course of the game.

From an abstract perspective, sensors can be regarded as generating values that correspond to various states, eg first half, second half etc. A 'profile' is a combination of various states. Each state occurs once and in the order specified. The goal is to semantically enrich sensor data with an additional field that identifies the state associated with every sensor reading. Our method is to convert the sensor stream to XML, which facilitates the subsequent semantic enrichment process. In simple terms, the synchronization process involves identifying one or more specific moments in time, such as the beginning or end of the game. Once the reading corresponding to that time is identified, the data can be synchronized with the data from all the other devices involved in the experiment.

The sensors used in our experiments record a heart rate value every 5 seconds, and approximately 1200 values are generated while the device is worn. The six states corresponding to a Gaelic football match can be seen in Figure 1. This example graph is for a midfield player who has a profile of gradually increasing activity through Pre-Game and Warm-Up, and remaining constantly active throughout each half. This profile can be easily split into states because of the period of rest located between the two periods of high activity. However, this profile is atypical among the thirty players involved in a given game. A more typical graph is shown in Figure 2, corresponding to a defensive player. This graph is characterized by short bursts of activity interspersed with periods of rest, making it much more difficult to correctly identify the state boundaries. This provides a significant challenge to creating a generic process for normalizing and synchronizing sensor streams.

In order to perform our synchronization, we define a 'model' profile of the ideal shape of the data graph. This comprises two periods of consistently high activity on either side of a period of relatively low activity. This model profile is compared to the data from each sensor device until the closest match is found, in terms of intensity and dura-

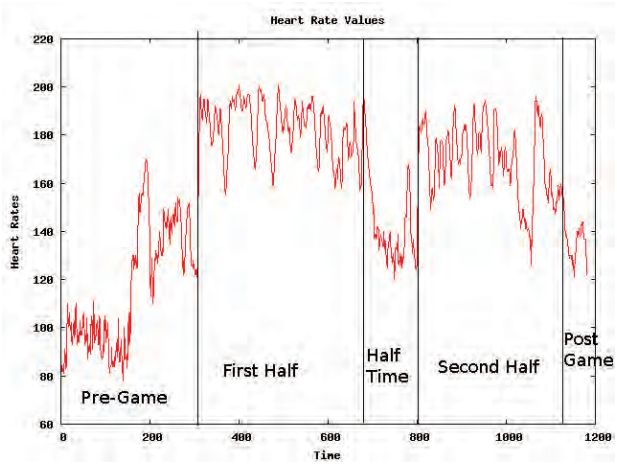


*Figure 1: Midfielder data with distinct state boundaries.*

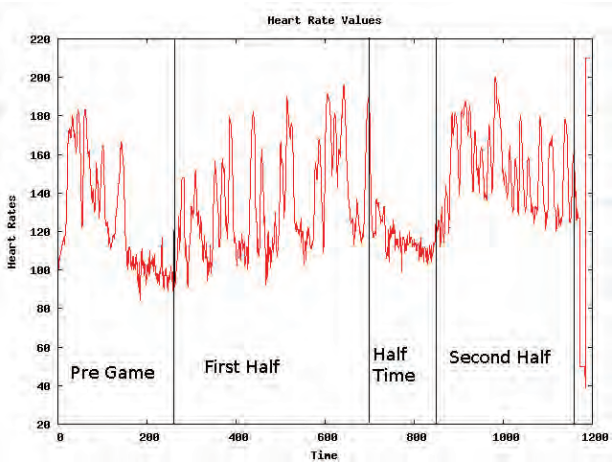


*Figure 2: Defensive player data with blurred state boundaries.*

tion of the activity. Identifying state boundaries for the resulting profile is reasonably simple, involving identifying the point at which the data changes in intensity and applying some rules from the domain experts to find the precise location on the curve at which the state changes. The durations of each state are then recorded based on these boundaries.

For the data from the other sensor devices, a single point is identified. In our experiments, we discovered that while in many cases it proved impossible to find precise state boundaries for all states, in every case it was possible to identify at least one, namely the beginning of the half-time period, when every player ceased intense activity and so their heart rates dropped significantly. Once this point is identified, the state durations extracted from the model profile, which should correspond to the durations of the game periods, are applied and each sensor reading is marked up with state information.

When all readings are assigned a state, it is possible to query the data according to state and across multiple sensors, resulting in richer knowledge and the potential for new knowledge discovery.

**Link:**
http://www.computing.dcu.ie/~isg

**Please contact:**
Dónall McCann, Mark Roantree, Niall Moyna, Michael Whelan
Dublin City University, Ireland
E-mail:
donall.mccann@computing.dcu.ie,
mark.roantree@computing.dcu.ie,
niall.moyna@dcu.ie,
michael.whelan@dcu.ie

# Inertial Sensing: A Little Bit of CLARITY

by John Barton, Brian Caulfield & Niall Moyna

***The increasing availability of cheap, robust and deployable sensor technology will usher in a new wave of ubiquitous information sources. A particular implementation of ambient sensors is in the area of wearable electronics in body area networks incorporating inertial sensing devices. As part of the CLARITY Centre for Sensor Web Technologies, the Tyndall Wireless Inertial Measurement Unit (WIMU) is being used in a number of projects focussing on two key themes: Health and Fitness, and Helping the Aged.***

## Wireless Inertial Measurement Unit

The Tyndall Wireless Inertial Measurement Unit (WIMU) is a 6 Degrees of Freedom (6DOF) inertial sensing device, comprising triple-axis accelerometers, gyroscopes (angular velocity) and magnetometers. The triple-axis acceleration and angular velocity sensor output values can be combined in a nonlinear matrix equation to give both position and orientation information. The system can be visualized by using a fixed frame of reference for position measurement (x, y, z), the Earth-Fixed Frame, and utilizing a moving non-inertial frame (u, v, w), the IMU-Fixed Frame, which has its axes parallel to those of the IMU sensors.

The 25mm WIMU was developed based upon Tyndall's 25mm modular wireless sensor node technology. It is one of a large family of layers currently available for the Tyndall25. The 25mm wireless node has been used to develop a platform for low-volume prototyping and research in the wireless sensor network domain. A number of research projects currently underway at the institute are using it as a platform for sensing and actuating in scalable, reconfigurable distributed autonomous sensing networks, and it is supported by Science Foundation Ireland (SFI) through Tyndall's National Access Program (NAP).

# CLARITY - Centre for Sensor Web Technologies

The CLARITY Cenre for Sensor Web Technology focuses on the intersection between two important research areas – adaptive sensing and information discovery – to develop innovative technologies of critical importance to Ireland's future and to improving quality of life in areas such as personal health, digital media and environmental management. The theme of CLARITY's research programme – bringing information to life – refers to the harvesting and harnessing of large volumes of sensed information, both from the physical world in which we live and the digital world of modern communications and computing. CLARITY commenced in June 2008 and represents a large-scale academia-industry collaboration, accommodating more than 100 full-time researchers from University College Dublin, Dublin City University and the Tyndall National Institute, in partnership with more than ten industrial partners, including major multinationals and emerging Irish companies.

## Inertial Sensing for Health and Fitness

The development of unobtrusive sensing elements embedded in the fabric of garments has opened countless possibilities for the innocuous monitoring of athletes over extended periods of time in a variety of sport settings. Foster Miller (an independent company but part of the QinetiQ Inc. group) has recently developed a T-shirt-based Ambulatory Physiological Monitoring System, which monitors the vital signs of a person during activity and transmits the data wirelessly to a remote station.

The combination of textile sensors with WIMUs will greatly assist in the ambulatory monitoring of healthy individuals and of those with chronic diseases such as obesity, diabetes, heart failure, and arthritis. The information will allow patients and allied health professionals to monitor physiological response during various forms of activity, and to design individually tailored programs for individuals across their lifespan. In addition, the technology has the potential to allow sports coaches and trainers to monitor individual athletes in a training or competitive environment. Currently, coaches/trainers are very limited in what they can measure in real time in a training or competitive environment. The information from the proposed sensor platform could be used to design training programs that replicate the movement patterns and/or physiological responses for a given sport or for a specific position (eg fullback vs. centre forward) in a team sport.

For biomechanical analysis, the ability to monitor athletes' movement in their natural environment is a huge leap forward compared to the current method of measuring them in a laboratory setting. As part of CLARITY, critical markers such as the speed and agility of top-level rugby players will be determined in the lab. They will then be outfitted with WIMUs for assessment in the field, to examine for example the 'cutting' movement of players when they make a rapid change in direction and the stresses and strains that these intense motions put on their joints and muscles.

As part of a collaboration with Tennis Ireland, tennis players will be fitted with WIMUs to augment the already rich sensory environment available at the tennis facilities at Dublin City University. An all-weather tennis court has been instrumented with nine Internet-enabled cameras with built-in microphones. Adding WIMUs to the tennis players' bodies will enable us to determine the actions they are performing and even the stroke they are playing.

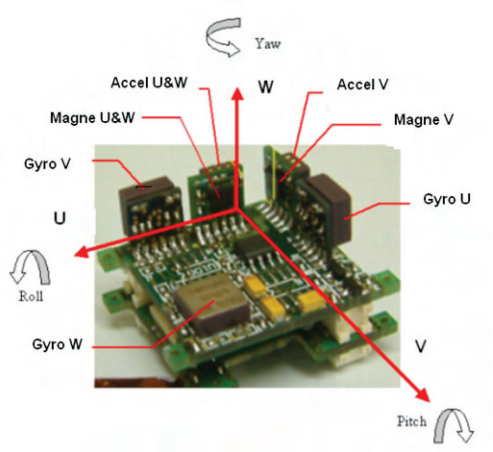


*Figure 1: Tyndall Wireless Inertial Measurement Unit showing yaw, pitch and roll.*

**Links:**

http://www.clarity-centre.com/
http://www.tyndall.ie/mai/25mm.htm
http://www.tyndall.ie/nap/
http://www.dcu.ie/shhp/index.shtml
http://www.ucd.ie/physioperformsci/
http://www.foster-miller.com/
http://www.cdvp.dcu.ie/tennisireland/

**Please contact:**

John Barton
Tyndall National Institute, Cork, Ireland
Tel: +353 21 4904088
E-mail: john.barton@tyndall.ie

Brian Caulfield,
University College Dublin, Ireland
Tel: +353 1 7166502
E-mail: b.caulfield@ucd.ie

Niall Moyna,
Dublin City University, Ireland
Tel: +353 1 7008802
E-mail: niall.moyna@dcu.ie

# Semantically Enhanced Representation of Legal Contracts for Web Applications

by Mihály Héder and Balázs Rátai

***Carneades Contract Format (CCF) is a flexible and extensible representation framework for legal contracts. This new format makes it possible to represent equally the text, document structure and semantics of legally binding agreements. The concept was developed by Carneades Consulting and the Internet Technology Department of SZTAKI, which also played an important role in the selection of the optimal technical solution and developed the first demo implementation.***

The majority of Web applications require the user's consent to various types of contract before access to their services can be granted. A typical example of such a contract is a 'terms of use' document. Online commercial transactions also result in contractual relationships (eg sales contracts in the case of Web shops).

The current practice of online contract management has many shortcomings, the most significant of which are the following: (i) agreements appear in the form of long and complex text documents that cannot be processed with machines; (ii) overwhelmingly general contractual terms and conditions are used; and (iii) the transaction-specific contractual information is stored separately in databases.

One of the reasons for this situation is the lack of a generic standard for the representation of legally binding agreements that is simple and flexible enough to store the text of the agreement, the document structure and the related semantic data in the same document. The CCF intends to be just such a format. Instead of producing a rich semantic expression set or a data dictionary for a specific field, we have been focusing on the structure of the document. Having examined the everyday practice of legal procedures and the requirements of lawyers, we found that there is a strong need for electronic legal documents which are usable as documentary evidence in legal procedures. For this to be possible, all the relevant information should be stored in the same place (ie not distributed in different databases, transaction logs and legal texts about the general conditions of use), should be in a format readable by lay users (so that experts are not needed to interpret it), and the document should be provably authentic (ie it should be signed). Ideally, documentary evidence should be very easy to present to the court. Our aim is to develop a format specifically to meet these criteria. To the best of our knowledge, this is a completely new approach to developing a legal format.

In order to help to overcome the above-mentioned shortcomings of present practice, we have defined the following design goals:

- the format must be as simple as possible
- the format should capture only the structure of a contract document, allowing the user to choose the semantic vocabulary (ie ODRL Data dictionary elements or XACML expressions) of the annotation, thus allowing a wide range of existing legal semantic markup formats to be used
- the documents in this format must be readable with a common Web browser
- it should be possible to embed the documents using the format into existing Web pages in order to allow the development of interoperable Web solutions.

In order to meet the above requirements, a relatively simple XML schema was defined that describes the entire contract as a hierarchy of conditions and sub-conditions, references and annexes. The result is a CCF document that is capable of structuring the text of the contract into a simple, hierarchical form. It allows further annotation of the text with existing, sophisticated, application-specific formats such as Open Digital Rights Language (ODRL) or eXtensible Access Control Markup Language (XACML).

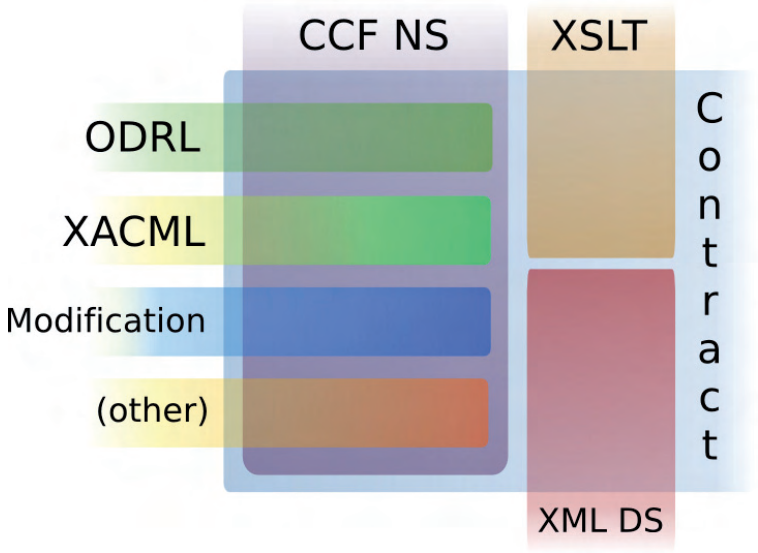


*Figure 1: The interplay of different namespaces in a contract.*

The next step is to embed the whole document into another XML file, which contains an Extensible Stylesheet Language Transformation (XSLT) document. The result is a document that, when opened by a Web browser, is a compound structure of XHTML+CSS+Javascript elements.

This kind of bundling of data and representation is a novel solution for legal contracts. The flexible structure of the document also allows an enveloped XML Digital Signature to be added.

A document created this way unifies the advantages of accessible Web documents and the expression power of the application-specific legal formats, which are capable of representing legally relevant and easily processed semantic information about the text.

As mentioned above, it is possible to embed, anywhere in the document, valid XML from any namespace (eg ODRL or XACML) in so-called extension elements.

By default, the semantic annotations in the document become simple visual annotations in the browser. If there are parts of the XSLT prepared to handle the particular namespace used in an extension element, the visual representation becomes even more informative.

A special kind of annotation is the contract modification, which we have placed in a different namespace. Using the elements of this namespace, we can produce contracts which modify other contracts. Using the original contract and the modifications we can always derive the current, consolidated view of the original and the modifying contract. On a technical level, however, the modification schema is just like an external schema to the contract, such as XACML, ODRL or others.


*Figure 2: An ODRL-annotated usage agreement document about some pictures bought in a Web shop. When opened in a browser, moving the cursor over the annotated (light red) text causes a pop-up containing the details.*

There are many possible applications of the format. Just to mention a few, it is possible to (i) derive access control decisions directly from XACML-annotated usage agreements; (ii) provide a summary of an annotated document; (iii) implement smart searching in documents based on semantics; and (iv) provide an overview of many contracts or generate statistics based on large number of annotated documents.

We are currently testing and evaluating our solution in the framework of a research project called AAI-Based Authorization Broker, an e-commerce solution based on strong identity and policy management. This is a sub-project of the Mobile Innovation Centre, a research and development program sponsored by the Hungarian government, in the area of mobile telecommunication.

**Links:**
AAI project:
http://www.sztaki.hu/search/projects/project_information/?uid=00210
CCF schema:
http://www.carneades.hu/xml/carneadescontract

**Please contact:**
Mihály Héder
SZTAKI, Hungary
Tel: +36 1 279 6027
E-mail: mihaly.heder@sztaki.hu

Balázs Rátai
Carneades Consulting, Hungary
Tel: +36 1 394 2114
E-mail: balazs.ratai@carneades.hu

# Contiki: Bringing IP to Sensor Networks

by Adam Dunkels

***The open-source Contiki operating system brings IP, the Internet Protocol, to sensor networks through the uIP (micro Internet Protocol), uIPv6 protocol stacks and the SICSlowpan IPv6-over-802.15.4 adaptation layer.***

Contiki is an open-source, memory-efficient operating system for sensor network nodes that was the first operating system to provide IP connectivity for sensor networks. Contiki incorporates many recent research results in wireless sensor networks, such as power profiling, cross-layer simulation, and low-power radio networking. The Contiki project was started in 2002 and has subsequently grown to include research institutions and major industry players. At SICS, we use Contiki for several sensor network research projects.

The Contiki operating system is implemented in C and consists of an event-driven kernel, on top of which application programs can be dynamically loaded and unloaded at run time. Contiki processes use lightweight protothreads that provide a linear, thread-like programming style on top of the event-driven kernel. In addition to protothreads, Contiki also supports per-process optional preemptive multithreading and interprocess communication using message passing. Contiki runs comfortably in a few kilobytes of RAM.

Several of Contiki's mechanisms have been released as separate open-source packages and have seen significant industrial uptake. The uIP embedded IP stack, originally released in 2001, is today used by hundreds of companies in systems such as freighter ships, satellites and oil drilling equipment. Contiki's protothread programming library, first released in 2005, has been used in digital TV decoders and wireless vibration sensors.

For sensor network communication, Contiki provides a low-power radio net-

working stack called Rime. The Rime stack implements sensor network protocols ranging from reliable data collection and best-effort network flooding to multi-hop bulk data transfer and data dissemination. IP packets are tunnelled over multi-hop routing via the Rime stack.

Interaction with a network of Contiki sensors can be achieved with a Web browser, a text-based shell interface, or dedicated software that stores and displays collected sensor data. The text-based shell interface is inspired by the Unix command shell but provides special commands for sensor network interaction and sensing.

To provide a long sensor network lifetime, it is crucial to control and reduce the power consumption of each sensor node. Contiki provides a software-based power profiling mechanism that keeps track of the energy expenditure of each sensor node. Being software-based, the mechanism allows power profiling at the network scale without any additional hardware. Contiki's power profiling mechanism is used both as a research tool for experimental evaluation of sensor network protocols, and as a way to estimate the lifetime of a network of sensors.

Contiki provides a flash-based file system, called Coffee, for storing data inside the sensor network. The file system allows multiple files to coexist on the same physical on-board flash memory and has a performance that is close to the raw data throughput of the flash chip.

To ease software development and debugging, Contiki provides three simulation environments: the MSPsim emulator, the Cooja cross-layer network simulator, and the Netsim process-level simulator. The development process for software for Contiki typically goes through all three simulation stages before the software runs on the target hardware.

In October 2008, major industry players Cisco and Atmel joined Contiki. Cisco, Atmel and SICS jointly announced uIPv6, the world's smallest fully compliant IPv6 stack. uIPv6 builds on the uIP stack and is integrated in Contiki.

The Contiki team currently consists of sixteen developers from SICS, SAP AG, Cisco, Atmel, NewAE and TU Munich.

*Adam Dunkels from SICS, Sweden, is the winner of the 2008 Cor Baayen Award for a promising young researcher in computer science and applied mathematics.*
*http://www.ercim.org/activity/cor-baayen*

**Link**:
http://www.sics.se/contiki/

**Please contact:**
Adam Dunkels
SICS, Sweden
Tel: +46 70 773 1614
E-mail: adam@sics.se

# Preferential Text Classification: Learning Algorithms and Evaluation Measures

by Fabio Aiolli, Riccardo Cardin, Fabrizio Sebastiani and Alessandro Sperduti

***Researchers from ISTI-CNR, Pisa and from the Department of Pure and Applied Mathematics at the University of Padova, are explicitly attacking the document classification problem of distinguishing primary from secondary classes by using 'preferential learning' technology.***

In many contexts in which textual documents are labelled with thematic classes, a distinction is made between the primary and secondary classes to which a given document belongs. The primary classes of a document represent the topic(s) that are central to the document, or that the document is mainly about. The secondary classes instead represent topics that are somehow touched upon, albeit peripherally, and do not represent the main thrust of the document.

This distinction has been neglected in text classification (TC) research. We contend that it is important and deserves to be explicitly tackled since, in most contexts in which the distinction is made, the degree of importance of a misclassification can depend on whether it involves a primary or a secondary class. For instance, when a patent application is submitted to the European Patent Office (EPO), a primary class from the International Patent Classification (IPC) scheme is attached to the application, and that class determines the expert examiner who will be in charge of evaluating the application. Secondary classes are attached only for the purpose of identifying related prior art, since the appointed examiner will need to determine the novelty of the proposed invention against existing patents classified under either the primary or any of the secondary classes. Thus, for the purposes of the EPO, failing to recognize the true primary class of a document is a more serious mistake than failing to recognize a true secondary class. Similar considerations apply to other scenarios in which the distinction is made.

In a concerted attempt to address this distinction, we define preferential text classification, a task which we define as the attribution to a textual document d of a partial ordering among the set of classes C. This partial ordering specifies whether or not a given class 'applies more than' (or 'is preferred to') another class in the document. In particular, we focus on a special case of preferential TC; namely, the case in which each document is associated to a 'three-layered' partial order. This consists of a top layer of one or more primary classes, each of which is preferred to those in a middle layer of secondary classes, which are in turn each preferred to those in a bottom layer of 'non-classes' (ie classes that do not apply at all to the document).

The original contribution of our work is twofold. First, we propose an evalua-

tion measure for preferential TC, in which different kinds of misclassifications involving either primary or secondary classes have a different impact on effectiveness. Second, we attack preferential TC by using a learning model, dubbed the Generalized Preference Learning Model, that was explicitly devised for learning from training data expressed in preferential form, ie in the form "class c' is preferred to class c" for document d". This model allows us to draw a fine distinction between primary and secondary classes in both the testing and learning phases, thus making use of the different importance of primary and secondary classes to which a training document belongs. Experiments run on WIPO-alpha, a well-known benchmark dataset consisting of manually classified patents, show that the Generalized Preference Learning Model outperforms standard (ie non-preferential) state-of-the-art learning approaches.

**Link:**
http://www.isti.cnr.it/People/F.Sebastiani/Publications/IRJ08b.pdf

**Please contact:**
Fabrizio Sebastiani
ISTI-CNR, Italy
Tel: +39 050 3152 892
E-mail: fabrizio.sebastiani@isti.cnr.it

# DataCell: Exploiting the Power of Relational Databases for Efficient Stream Processing

by Erietta Liarou and Martin Kersten

***Designed for complex event processing, DataCell is a research prototype database system in the area of sensor stream systems. Under development at CWI, it belongs to the MonetDB database system family. CWI researchers innovatively built a stream engine directly on top of a database kernel, thus exploiting and merging technologies from the stream world and the rich area of database literature. The results are very promising.***

Rather than simply transmitting the raw measured data, current state-of-the-art sensors are capable of a limited amount of processing. This feature has many positive effects, such as keeping the network usage and costs as low as possible. However, this is not enough to replace the role of well-equipped nodes that gather streaming sensor data from multiple sources and which account for the biggest share of the processing cost. These nodes should be able to perform complex query processing on large amounts of incoming data, meeting strict real-time deadlines even in periods when the frequency of incoming data explodes.

Our work focuses on this part of the sensor research. We are designing and developing a system called the DataCell, which is capable of efficiently collecting and processing high volumes of stream data. We are currently studying the DataCell over the stream application scenario of an ambient home setting. The DataCell is positioned as a data refinery cell that acts as an easily programmable data hub in a multi-network environment. Its task is to collect, filter and aggregate information from different sources to enable complex decision making, such as control of the lighting based on audio/video presentations. The challenge in an ambient environment is to hide the computer from the casual user, even while it is actively steering the environment. An example query in the ambient scenario could be, "tune the

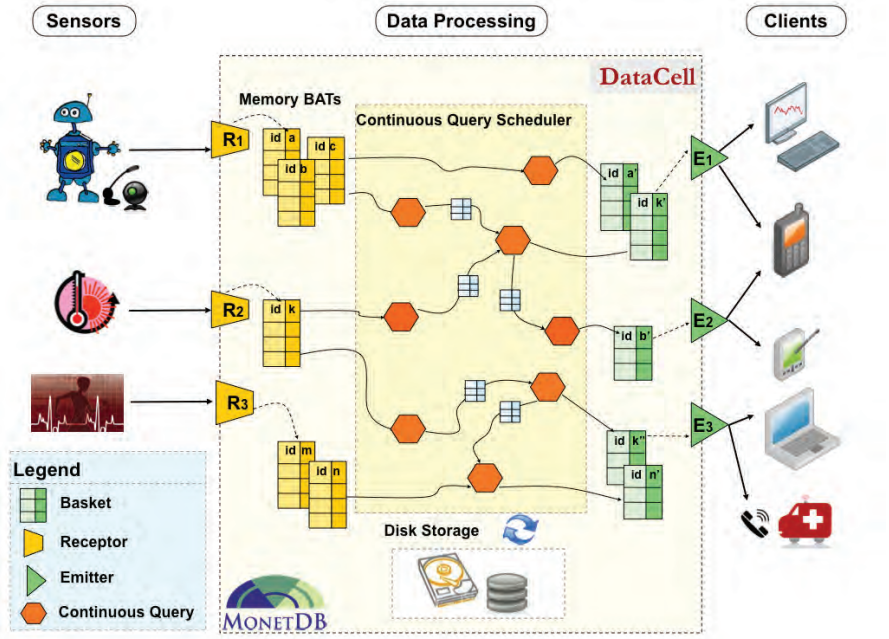


*Figure 1: The DataCell in the ambient scenario.*

television to my favourite show when I sit on the couch"; ie depending on the weight measured by a sensor in the seat, and the time of day, different TV shows will appear on the screen.

In stream applications, we need mechanisms to support long-standing queries over data that is continuously updated from the environment. This requirement is significantly different from what happens in a traditional database system, where data are stored in static tables and users fire one-time queries to be evaluated over the existing data. Given this critical difference, the pioneering architects of the data stream management system naturally considered existing database architectures inadequate to achieve the desired performance: instead they designed new architectures from scratch.

However, working from scratch makes it difficult to exploit the existing knowledge and techniques of relational databases. This disadvantage became more pronounced as the stream applications demanded more functionality. In DataCell therefore, we started at the other end of the spectrum, building an efficient data stream management system on top of an extensible database kernel. With careful design, this allows us to reuse the sophisticated algorithms and techniques of traditional databases. We can provide support for any kind of complex functionality without having to reinvent solutions and algorithms for problems and cases for which a rich database literature already exists. Furthermore, it allows for more flexible and efficient query processing by allowing batch processing of stream tuples, as well as non-consecutive processing by selectively picking the tuples to process.

The idea is that when stream tuples arrive in the system, they are immediately stored in (appended to) a new kind of table called a basket. By collecting tuples into baskets, we can evaluate the continuous queries (which are already submitted to the system and are waiting for future incoming data) over related baskets as if they were normal one-time queries. This allows us to reuse any kind of algorithm and optimization designed for a modern database system. Each query has at least one input and one output basket. It continuously reads data from the input baskets, processes this data and creates a result which it then places in its output baskets. Once a tuple has been seen by all relevant queries, it is dropped from its basket.

This description of the process is somewhat simplified, since this process allows the exploration of quite flexible strategies. For example, the same tuple may be thrown into multiple baskets where multiple queries are waiting, query plans may be split into parts, and baskets may be shared between similar operators (or groups of operators) of different queries, allowing results to be reused.

The periphery of a sensor stream engine is formed by adapters, eg software components to interact with devices, RSS feeds and SOAP Web services. The communication protocols range from simple messages to complex XML documents transported using either UDP or TCP/IP. The adapters for the DataCell consist of receptors and emitters. A receptor is a separate thread that continuously picks up incoming events from a communication channel and forwards them to the DataCell kernel for processing. Likewise, an emitter is a separate thread that picks up events prepared by the DataCell kernel and delivers them to interested clients, ie those that have subscribed to a query result.

We designed and developed the DataCell at CWI in Amsterdam, funded by the BRICKS project. It is implemented on top of the MonetDB, an open-source column-oriented database system. Currently it is a research prototype and the goal is to be able to disseminate the DataCell soon as part of MonetDB.

**Link:**
http://monetdb.cwi.nl/

**Please contact:**
Erietta Liarou
CWI, The Netherlands
Tel: +31 20 59 24 127
E-mail: erietta@cwi.nl

Martin Kersten
CWI, The Netherlands
Tel: +31 20 59 24 066
E-mail: mk@cwi.nl

# On looking FORWARD

by Sotiris Ioannidis, Evangelos Markatos and Christopher Kruegel

***Computer systems, networks and Internet users are under constant threat from cyber attacks. FORWARD is an initiative by the European Commission to promote collaboration and partnership between academia and industry in their common goal of protecting Information and Communication Technology infrastructures.***

The past few years have been marked by an ever-increasing number of cyber attacks. Motivated by fun, fame and peer recognition, early attackers, more widely known as 'hackers', pioneered the methods used to penetrate computers, compromise accounts and invade our personal lives. Even though these early hackers usually meant no harm, their methods and techniques perfected the necessary technology required to compromise remote computers. In turn, this paved the way for professional criminals motivated by profit to start using compromised computers for a wide variety of illegal activities, such as trading of credit card numbers, online renting of compromised computers, online ordering and delivering of denial-of-service attacks, and sending spam email messages. To reduce the effects of these cyber attacks, security researchers are engaged in an arms race against the ever-increasing sophistication of cyber attackers, by creating systems that detect, and whenever possible mitigate, the effects of these attacks.

To stay ahead in this arms race, FORWARD brings together European

researchers in network and information systems security to identify (i) the most probable security threats in the near future, and (ii) those research areas that must be pursued to address and mitigate these emerging threats. By mobilizing a critical mass of researchers in Europe and by complementing them with a select team of researchers from Asia and America, FORWARD is working towards establishing a research agenda for cyber security in Europe and identifying possible new areas and threats that must be addressed. FORWARD researchers have focused their activities on three critical domains:

- Malware and Fraud: malware is perhaps the one arena in which attackers have clearly demonstrated an increased sophistication. In its race to evade antivirus signatures and systems and stay below the detection 'radar', malware has evolved to be agile, stealthy and highly sophisticated.
- Smart Environments: the increasing miniaturization of computing systems is driving the penetration of intelligent appliances in every human activity. As computing and communicating devices become increasingly widespread, so does the potential of attackers to disrupt our daily lives in a wide variety of ways.
- Critical Systems: Our daily functions, if not our lives, depend on a wide variety of traditional and emerging infrastructures, such as the power grid and communications networks. As it becomes more common to connect critical infrastructures to the Internet using off-the-shelf technologies, the vulnerability of these utilities increases to breaches and attacks from the outside world.

By mobilizing cyber security researchers in Europe and by consolidating their efforts along those major research axes, FORWARD will identify those research directions that will help lead to a safer and more secure cyberspace for all European citizens.

For more information about the activities of FORWARD or if you are interested in participating, please contact Christopher Kruegel or visit our Web site.

**Link:**
http://www.ict-forward.eu

**Please contact:**
Christopher Kruegel
Vienna University of Technology, Austria
E-mail: chris@seclab.tuwien.ac.at

Sotiris Ioannidis
FORTH-ICS, Greece
Tel: +30 2810391945
E-mail: sotiris@ics.forth.gr

# Enhancing Authentication in eBanking with NFC-Enabled Mobile Phones

by Diego A. Ortiz-Yepes

***In the past few months, a mobile phone-based authentication mechanism for eBanking has been developed at the IBM Zurich Research Laboratory. At the core of this mechanism, we have used NFC and CAP. The latter, Chip Authentication Program (CAP), is a specification developed by MasterCard that provides mechanisms for customer authentication based on smart cards compliant with EMV (Europay - MasterCard - Visa). The former, Near-Field Communication (NFC), is an emerging technology related to RFID that is already being incorporated into commercially available mobile phones, allowing them to communicate over very short distances (in the order of a few centimetres) with other NFC-enabled devices. This ability, when employed in tandem with CAP — as we have done in our authentication mechanism — greatly enhances the overall usability of the authentication system.***

Our NFC-based authentication mechanism relies on dual-interface smart cards, that is, cards with both contact and contactless interfaces. These cards might also be used for other financially related purposes, eg as debit or credit cards. In fact, this situation is desirable in order to avoid burdening the customer with an additional card for eBanking purposes.

The customer authentication mechanism works by having the customer produce an appropriate response to an unpredictable challenge generated by the bank. In order to do so, she must use her card and its PIN, which is used to authenticate the customer to the card. More precisely, when the customer wishes to engage in eBanking, she visits the Internet site of her bank, which requests her customer ID, eg her account or contract number. Once such an ID has been received by the bank, it replies with a challenge, which consists of an unpredictable number of between 6 and 8 digits. Having received this challenge, the customer starts the phone application by touching her bank card to the back of the phone (see Figure 1). She then selects the log-in mode and types in the server-issued challenge. Prior to generating the corresponding response, the phone requests that the customer provide her PIN in order to authenticate herself to the card. Once the customer has been authenticated by the card, the phone sends the challenge to the card obtaining a cryptogram in return. Using this cryptogram – a bit-string cryptographically bound to the challenge and the internal card state – the phone generates a numeric code, ie the response, which is displayed to the customer. Subsequently, she sends the response to the bank server by typing it into the PC. When the response is received by the bank, the latter checks whether it corresponds to the previously issued challenge. If this is the case, the bank presents the customer with her account(s) summary, as well as some appropriate transaction options.

The mechanism outlined above replaces the Personal Card Reader (PCR) required by some authentication schemes currently in use, yielding a


*Figure 1: Using NFC-enabled phones for eBanking authentication.*

more convenient authentication mechanism. This follows from the fact that the user needs only her phone and her card in order to authenticate herself to the bank. On the one hand, phones are truly ubiquitous devices that can hardly be considered a burden; on the other, most people carry their bank cards with them in their wallets or purses, making the requirements of our authentication mechanism quite low.

Note that both the challenge and the response could be sent directly from the bank server to the phone and back via SMS, or some other suitable mechanism using the mobile phone network. This would not only simplify the mechanism, but would also increase the level of security as a consequence of using the phone and a secondary channel, whose compromise is much less likely than the PC alone.

In conclusion, we have developed an eBanking authentication mechanism whose security properties are comparable to PCR-based CAP. NFC-enabled mobile phones are a key component of this mechanism, providing an enhanced level of usability, not only by replacing the PCR altogether, but also being able to offer a more pleasant user experience. Additionally, we have implemented this mechanism in such a way that it integrates seamlessly with an existing CAP infrastructure, allowing it to be used in a real-life pilot by the end of 2008.

**Please contact:**
Michael Baentsch, Michael Osborne, Diego A. Ortiz-Yepes
IBM Zurich Research Lab, Switzerland
E-mail:
mib@zurich.ibm.com,
osb@zurich.ibm.com,
ort@zurich.ibm.com

# Sino-French IT Lab in Beijing Opens to European Institutions

by Stéphane Grumbach

*After twelve years, the Sino-French IT Lab in Beijing has opened its doors to European partners. Created in 1997 by INRIA and the Chinese Academy of Sciences to promote cooperation between France and China in IT, LIAMA ('Laboratoire d'Informatique, d'Automatique et de Mathématiques Appliquées') has conducted more than 100 research projects associating laboratories in the two countries.*

Last October, the organization of LIAMA was revisited in order to strengthen its impact on cooperation between Europe and China. A consortium was created to encourage interaction between institutions sharing the same vision in ICT, a field which is a top research priority for both Europe and China.

On the Chinese side, Tsinghua University together with the three institutes of the Chinese Academy of Sciences specializing in IT in Beijing, the Institute of Automation, the Institute of Software, and the Institute of Computing Technology, joined the consortium. All these institutions host research teams with European partners working on joint research projects. LIAMA also opened its consortium to companies active in R&D, to develop links with the industry of both regions. France Telecom was the first company to join the consortium, and discussions are taking place with corporations and SMEs with which LIAMA has cooperative projects.

LIAMA is also a hub for European students to study in China or undertake internships in the framework of their European curriculum. Several dozen foreign students work every year in LIAMA in close partnership with their Chinese counterparts, giving them a unique opportunity to understand better the culture and start learning the language. A reasonable proportion of the European students who work in LIAMA choose to stay or return later to China to work.

Of course, the historical partners of LIAMA play a fundamental role in the consortium, the first steering committee of which met last October. LIAMA was created by Professor Ma Songde, who studied and worked in France in the 1980s and became director of the Institute of Automation after he returned to China. The Institute of Automation has been a major partner in LIAMA's projects, and is hosting the LIAMA office for the consortium.

In June 2008, LIAMA also became a National Centre for International Research, a newly created structure supported by the Ministry of Science and Technology, to foster international cooperation. Since its creation. LIAMA has enjoyed sustained support from the China's authorities and research institutions.

The strength of LIAMA is a result of the presence of a large group of researchers and students from France staying for long periods. Currently, seven French researchers from INRIA and CNRS are working in LIAMA projects for terms of several years. Among these, four are working in the School for Information Science and Technology of Tsinghua University. These sustained relations with our Chinese partners have allowed the development of an atmosphere of confidence and trust, in which the geographic distance and the cultural differences between the two regions vanish. The Institute of Automation was the first in the Chinese Academy of Sciences to hire a permanent foreign researcher, and to award the habilitation to a foreigner, which demonstrates its strong trust and involvement in this partnership.

*Tomato profile from calibrated GreenLab model. GreenLab is a stochastic, functional and interactive model for plant growth, one of LIAMA's research topics.*

Since LIAMA started as a spin-off from the National Laboratory for Pattern Recognition of the Institute of Automation, image analysis has been one of the important directions of research. Over the years, our spectrum of activity has broadened. Currently, LIAMA's research is structured around the following disciplines:

- computational medicine, with a special focus on brain imaging
- environmental modelling, from the modelling of plant growth to the study of the turbulence of sand winds
- secure software, formal methods for embedded software, network and grid programming, as well as open-source software promotion
- graphics, computer-aided design and interactive technology
- pattern recognition and image understanding and processing.

LIAMA has contributed to establishing a privileged partnership between France and China for research in ICT. The LIAMA consortium will deepen our cooperation with the best partners in Europe and China.

European students can check for internship offers online, or directly contact the researchers in their field of interest to find out more about available positions. Institutions or labs can contact the directors of LIAMA for further information.

**Link:** http://liama.ia.ac.cn/

**Please contact:**
Stéphane Grumbach
INRIA, France, Director of LIAMA
E-mail: stephane.grumbach@inria.fr

# DocSearch: UHF RFID Technology for Document Localization

by Marina Buzzi, Marco Conti and Daniele Vannozzi

***Radio-frequency identification (RFID) and sensor technologies are now enabling the Internet of Things. Ultra High Frequence (UHF) RFID readers and passive tags can be valuable tools for emerging pervasive services since they allow tags to be read at distances ranging from one-half to few meters, depending on antenna power, size and polarization. However, UHF has limits due to RF reflection, shadowing and absorption. Our experimental study investigates the feasibility of UHF RFID for reliable and efficient retrieval of archived documents.***

Radio-frequency identification (RFID) uses radio waves to permit the automatic identification of objects, people and animals. It consists of two components, readers and tags; the tags store information that can be retrieved by readers. Tags may be passive, active (battery-powered) or semi-active. Passive tags are especially convenient since they are small, cheap and potentially can last infinitely.

The DocSearch project is underway in Pisa, Italy, at the Istituto di Informatica e Telematica of the Italian National Research Council. It aims to develop a tool for improving the efficiency of the ccTLD ".it" Registry which assigns and manages domain names under the country code Top Level Domain "it". Specifically, the goal of the project is to aid Registry staff to answer customer questions regarding paper documents submitted as part of a registration request. The operator needs an efficient method to retrieve the original document(s) referred to by the customer from among thousands of documents. The DocSearch project is thus developing a tool that exploits RFID technology. Each document is archived with an RFID tag and the document data and tag id are stored in a database. Using RFID technology the operator can easily find the correct document even if the document was stored in the wrong place.

When applying RFID technology, the reliability of the reading is crucial. Raw RFID data are large-volume streams characterized by duplicate, missed and ghost reads. Therefore, data filtering and aggregation are necessary in order to extract reliable data.

UHF tag reading can be challenging depending on object composition, packaging, and tag size and placement. It is facilitated by motion. Three issues affect UHF reading: reflection, shadowing and absorption:

- metal reflects RF (however an appropriate insulation between tag and object can improve reading)
- shadowing is present when several tags are placed very close to one another and their antennas mask each other, and this decreases the read rate
- liquids (such as water) absorb RF and hinder tag reading.

The DocSearch project began in 2008 with a feasibility study aimed at measuring the reliability of UHF tag reading applied to document search. In the feasibility study we tested both multiple reads and single search. Multiple simultaneous reads (usually called inventory mode/command) make it possible to discover multiple tags in the antenna field at the same time, by using an anti-collision search algorithm. In order to verify the amplitude of the shadowing problem, we tested the worst case when each document consists of a single paper sheet and all sheets are piled in a folder. We carried out experiments in several configurations with different spatial positioning of tags. The best result was achieved by applying a thin dielectric substrate, which separates the tag from the page (increasing distance between tags). In this configuration the mean read percentage was 85% (see Table 1).

Looking for a specific tag is less critical. The success percentage is almost 100% when searching a 50-document folder (see Table 2).

Since single document retrieval is the main operation in our scenario, the feasibility study indicates that UHF RFID is a suitable technology for efficient document localization; thus we have developed a procedure for the storage and retrieval of documents using RFID technology. Currently a preliminary prototype (developed in C# language) is available for testing on a WorkAbout Pro 2 palm (Psion Teklogix) integrating an RFID reader (CAEN/Intel technology), with Win Mobile 6 OS. Documents were tagged with ALN 9540 "Squiggle" (Alien Technology). The development environment is .NET, Visual Studio 2005 and SQL server 2005.

**Link:**
IT NIC http://www.nic.it/

**Please contact:**
Marina Buzzi, Marco Conti, Daniele Vannozzi, IIT-CNR, Italy
E-mail: marina.buzzi@iit.cnr.it, marco.conti@iit.cnr.it, daniele.vannozzi@iit.cnr.it

| | % success | Tag identified | Reads | Time in sec |
|---|---|---|---|---|
| Inventory | 85 | 34/40 | 710 | 60 |

*Table 1: Reading results in inventory mode.*

| | % success | Tag identified | Num tests | Time in sec |
|---|---|---|---|---|
| Single Search | 100 | 1/50 | 50 | 5-10 |

*Table 2: Reading results for the tag retrieval operation.*

# Integrated Site Security for Grids

by Kate Bradshaw

***Between 90 and 98% of the emails received each day by most organizations are spam. While some are nothing but a harmless nuisance, others are malicious and capable of causing substantial damage. This is just one of the ways in which Grid and other computing sites can be attacked. To increase awareness and provide security guidance, CERN has led a European Commission co-funded project entitled Integrated Site Security for Grids (ISSeG), which was completed at the end of March 2008. The final results of this 26-month project are available from the project Web site.***

The ISSeG Web site provides a risk assessment tool, security recommendations and training material to help sites improve their computer security. While the project's focus has been on the security of Grid sites, the material is applicable to a wide range of computer centres, particularly those in academic or technical environments.

## Integrated Site Security

The project's vision has been that Grid security, which focuses on middleware, authentication, authorisation and operation across multiple administrative domains, needs to be complemented by comprehensive site security at all participating Grid sites. To this end, the ISSeG project has created and disseminated practical expertise on the deployment of Integrated Site Security (ISS).

ISS is a practical approach to site security that integrates technical, administrative and educational security solutions, and develops them in a consistent and coordinated way. This integration ensures that policies, rules, awareness and training all evolve in step with technological or administrative developments.

## Creating Practical Expertise

The project began in February 2006 and has been co-funded by the EU FP6 Programme. The consortium comprised the European Organization for Particle Physics (CERN) in Switzerland, Forschungszentrum Karlsruhe (FZK) in Germany, and the Science and Technology Facilities Council (STFC) formerly known as the Council for the Central Laboratory of the Research Councils (CCLRC) in the UK.

ISSeG created and captured raw expertise through full-scale ISS deployment at CERN and FZK including, for example, flexible and improved security for centrally managed computers, strengthened policies for controls networks and increased firewall protections. Experience gained from the two site deployments, as well as site security assessments carried out by a subcontracted company, were used to develop training materials and recommendations as to how security risks can be mitigated.

Coordinated by STFC, results were disseminated via presentations and the ISSeG Web site to help scientific communities use this integrated approach to improve their site security measures. In addition, there has been close collaboration with the security groups of the Enabling Grids for E-sciencE (EGEE) project – the Operational Security Coordination Team (OSCT) and the Joint Security Policy Group (JSPG) – resulting in their continuous involvement and input to help shape ISSeG development.


*The ISSeG project complements Grid security by providing recommendations and training for each Grid site.*

## The ISSeG Web Site

Visitors to the Web site can download and complete a risk assessment questionnaire. They will then, via a prioritized list of threats specific to their site, receive tailored site security recommendations. A generic set of the top threats and top recommendations for Grid sites can also be viewed directly.

Training materials for general users, system administrators, software developers and managers are also available from the site, including introductory material, training presentations, security checklists and downloadable printable materials.

Acknowledging the usefulness of the content, the OSCT will take over the maintenance of the ISSeG Web site to ensure its continued availability beyond the lifetime of the ISSeG project. This agreement was formalized in a Memorandum of Understanding between ISSeG and phase three of the EGEE (Enabling Grids for E-science) project (EGEE-III), which began in May 2008.

This article is based on one published in the CERN Computer Newsletter.

**Link:**
http://www.isseg.eu

**Please contact:**
David Jackson
STFC, UK
E-mail: david.jackson@stfc.ac.uk

Book Review:

# "From computers to ubiquitous computing, by 2020"

ERCIM was requested by the Royal Society to review the *Philosophical Transactions of the Royal Society A: 366, 1881 pp 3663-3838: "From computers to ubiquitous computing, by 2020"; 28 October 2008; edited by Marta Kwiatkowska, Tom Rodden and Vladimiro Sassone.*

The publication collects papers from a discussion meeting held at the Royal Society in March 2008. The 18 individual papers each tackle an aspect of Ubiquitous Computing (UbiComp) ranging from RFID identifiers and global sensing, software and architecture to support UbiComp and embedded systems through to properties of systems including self-* properties, security and trust with additional short papers on healthcare, ethics, engineering, modelling and living with UbiComp. Each paper is of high quality but no real linking themes emerged; the publication exhibits very good work by researchers in a disconnected way. The introduction is more of a catalogue of what follows although in the early and late paragraphs the topic is introduced with application examples and some challenges are presented, respectively. The short papers on engineering (Crowcroft), modelling (Milner) and living in a ubiquitous world (Rodden) provide real broader context and insight but are too short to allow the ideas to be developed.

The publication provides much useful material for researchers in individual aspects of UbiComp and for certain aspects may well provide the base reference. For some aspects the work appears disconnected from parallel work where there are research results and a body of well-respected publications. Anyone wishing to obtain a high-level view of the significance of UbiComp, its importance in the future knowledge society, the challenges in its achievement with their relative importance and difficulty or the potential utility of the technology in various application areas is likely to be disappointed.

*by Keith Jeffery*

**More information:**
http://publishing.royalsociety.org/index.cfm?page=1851

Philosophical Transactions A is particularly interested in receiving unsolicited theme proposal For further information, please visit
http://publishing.royalsociety.org/philtransa/guest-editor/

# FMICS 2009 - 14th International ERCIM Workshop on Formal Methods for Industrial Critical Systems

Eindhoven, The Netherlands, 2-3 November 2009

The aim of the ERCIM FMICS workshop series is to provide a forum for researchers who are interested in the development and application of formal methods in industry. In particular, these workshops bring together scientists and engineers that are active in the area of formal methods and interested in exchanging their experiences in the industrial usage of these methods. These workshops also strive to promote research and development for the improvement of formal methods and tools for industrial applications.

Topics include, but are not restricted to:
- Design, specification, code generation and testing based on formal methods.
- Methods, techniques and tools to support automated analysis, certification, debugging, learning, optimization and transformation of complex, distributed, real-time systems and embedded systems.
- Verification and validation methods that address shortcomings of existing methods with respect to their industrial applicability (eg, scalability and usability issues).
- Tools for the development of formal design descriptions.
- Case studies and experience reports on industrial applications of formal methods, focusing on lessons learned or identification of new research directions.
- Impact of the adoption of formal methods on the development process and associated costs.
- Application of formal methods in standardization and industrial forums.

Publication of the workshop proceedings in the Springer series Lecture Notes in Computer Science (LNCS) is planned.

FMICS 2009 is part of the first Formal Methods Week (FMweek), which will bring together a selection of events in the area, including FM 2009 (16th symposium on Formal Methods), TESTCOM/FATES (conference on Testing of Communicating Systems and workshop on Formal Approaches to Testing of Software), PDMC (Parallel and Distributed Methods of verifiCation),FACS (Formal Aspects of Component Software), CPA (Communicating Process Architectures), FAST (Formal Aspects in Security and Trust), FMCO (Formal Methods for Components and Objects) and the REFINE Workshop.

**More information:**
http://2009.ecoop.org/
http://www.win.tue.nl/fmweek
http://www.inrialpes.fr/vasy/fmics/

# ECDL 2009 - 13th European Conference on Digital Libraries

Corfu, Greece, 27 September - 2 October 2009

The 13th European Conference on Digital Libraries (ECDL 2009) is organized by the Laboratory on Digital Libraries and Electronic Publishing, Department of Archives and Library Sciences, Ionian University and will be held from 27 September to 2 October 2009, on the island of Corfu, Greece.

ECDL is the major European forum focusing on digital libraries and associated technical, practical, and social issues. In an intense environment of transformations in digital libraries, ECDL 2009, under the general title "Digital Societies", invites submissions in the following categories: Full and Short Papers, Posters and Demonstrations, Workshops and Tutorials, Panels and Doctoral Consortium. The proceedings will be published by Springer, in the Lecture Notes in Computer Science Series. Papers of the Doctoral consortium will be published in the Bulletin of the IEEE-TCDL.

## Topics

Conference topics include, but are not limited to:

- *Infrastructures* (Digital Library Architectures; Technology for Digital Libraries Infrastructures (Grids, etc.); Interoperability; Generic Strategic Infrastructures; Domain Focused Infrastructures)
- *Content Management* (Metadata Schemas; Semi Structured Data; Data Interoperability and Integration; Digital Curation, Archiving and Preservation; Collection Development, Management, Policies and Legal Issues; Semantic Web Issues in Digital Libraries)
- *Services* (Information Retrieval; Multilingual and Multimedia information retrieval; Personalization in Digital Libraries; Ontologies and Knowledge Organization Systems; Social Networking & Web 2.0 Technologies; Log Data in Digital Libraries; User Interfaces)
- *Foundations (*Formal Issues in Digital Libraries; Conceptual Views of Digital Libraries; Legal Issues in Digital Libraries; User Studies & Digital Library Evaluation).

## Important Dates

*Papers, Posters and Demonstrations*
Submission deadline: 21March 21, 2009
Acceptance notifications: 11 May 2009
Camera ready versions: 31 May 2009

*Doctoral Consortium*
Submission deadline: 1 June 2009
Acceptance notifications: 30 June 2009

*Workshops, Tutorials and Panels*
Submission deadline: 27 February 2009
Acceptance notification: 15 April 2009

**More information:**
http://www.ecdl2009.eu/



Sponsored by ERCIM

# ETAPS 2009 - European Joint Conferences on Theory and Practice of Software

York, UK, 22-29 March 2009

ETAPS is the primary European forum for academic and industrial researchers working on topics relating to Software Science. ETAPS is composed of five main conferences:

- CC 2009, International Conference on Compiler Construction
- ESOP 2009, European Symposium on Programming
- FASE 2009, Fundamental Approaches to Software Engineering
- FOSSACS 2009, Foundations of Software Science and Computation Structures
- TACAS 2009, Tools and Algorithms for the Construction and Analysis of Systems.

**More information:**
http://www.cs.york.ac.uk/etaps09/

2nd Edition of the ERCIM-ETSI Infinity Initiative

# Bio ICT - The Heart in the Computer

Sophia Antipolis, France, 2-3 April 2009

Systems Biology, the new science of complexity of living systems, will herald a fundamental change in paradigm with the appearance of life models and simulation systems making possible computer assisted experiments for biological, medical and pharmaceutical research. As such, the modelling and simulation of life is one of the most promising research areas with great hopes for humanity.

Building on the success of last year's seminar on Ambient Computing, ERCIM and ETSI will welcome recognized experts who will share their vision of modelling and simulation of life while demonstrating how research and standards, supported by the European Commission, can benefit society and business. Delegates will include scientists, industry leaders, key EC decision-makers and major players in the field.

**More information:**
http://www.etsi.org/WebSite/NewsandEvents/2008_InfinityInitiativeTeaser.aspx



Sponsored by ERCIM

# 23rd European Conference on Object-Oriented Programming

Genova, Italy, 6-10 July 2009

The European Conference on Object-Oriented Programming will be held in July of 2009 at Palazzo Ducale in Genova. ECOOP 2009 is the 23rd edition of the ECOOP Conference Series, which aims at bringing together people, industrial and academic, interested in a wide spectrum of technologies related to object-oriented areas. The ECOOP conference lasts for a week. Plenary sessions take place the last three days of the conference, whereas monday and tuesday are dedicated to workshops.

ECOOP 2009 wishes to embrace a broad range of topics, including:

- analysis, design methods and design patterns
- concurrent, real-time or parallel systems
- databases, persistence and transactions
- distributed and mobile systems
- frameworks, product lines and software architectures
- language design and implementation
- testing and metrics
- programming environments and tools
- theoretical foundations, type systems, formal methods
- versioning, compatibility, software evolution
- aspects, Components, Modularity, Reflection
- collaboration, Workflow.

ECOOP 2009 will also host a Summer School consisting of prestigious tutorials on current topics in software, systems, and languages research. The scope of the ECOOP Summer School is the same as the conference itself: all areas relevant to object technology, including work that takes inspiration from or builds connections to areas not commonly considered object-oriented. Tutorials should introduce researchers to current research in an area, and/or to show important new tools that can be used in research.

**More information:**
http://2009.ecoop.org/

# ERCIM elects President and Vice-Presidents

During the Board of Directors meeting of ERCIM in Porto on 20 and 21 November, the current President of ERCIM, Keith G. Jeffery (STFC, UK) was re-elected as President for a further two-year term starting 1st January 2009.

Three new Vice-Presidents (previously there were two) were also elected for the same two-year term. We are pleased to announce the following appointments:

- Matthias Jarke (FhG, D), responsible for the External Relations Task Group, takes over from Arne Sølvberg (NTNU, Norway), who will retire mid-2009 and therefore did not stand for another term
- Simon Dobson (UCD, Ireland), responsible for the Structural Task Group, takes over from Jan Karel Lenstra (CWI, NL). Jan Karel did not wish to stand for re-election but will remain a member of this TG
- Staffan Truvé (SICS, Sweden), responsible for the newly founded Innovation Task Group, was elected as the third Vice-President.

# ERCIM meets Günter Verheugen

ERCIM, together with INRIA Transfert, a subsidiary of INRIA to create and develop IT companies, had a private meeting in Sophia Antipolis with Günter Verheugen, Vice-President of the European Commission and Commissioner

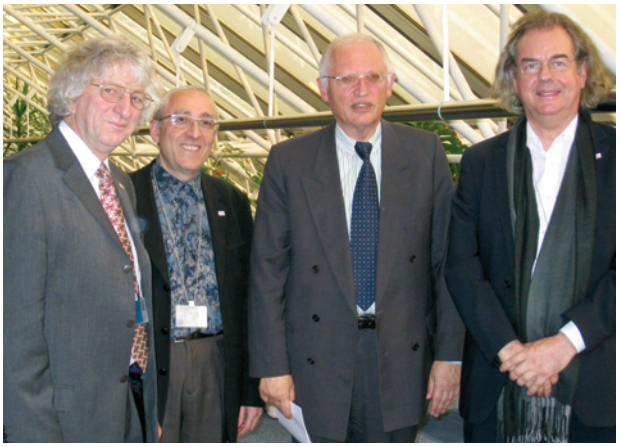

*From left: Keith Jeffery (ERCIM President), Laurent Kott (INRIA Transfert), Günter Verheugen and Jérôme Chailloux (ERCIM Manager).*

for Enterprise and Industry, during the conference 'Towards World Class Clusters' organised by the French Presidency of the European Union on 13 November. Commissioner Verheugen was very interested in ERCIM's positioning as a 'European network of research' becoming also a 'European network of innovation'. He invited ERCIM to participate in a working session on innovation to be organised by his Directorate General in Brussels for a small number of invited participants.

# STFC Staff celebrate Achievements of the Atlas Centre

In December 2008, the IT staff at STFC - who had been centred in the historic Atlas Building - held a small event to celebrate the achievements since the installation of UK's first supercomputer in 1964. In January the staff move to a new building labelled R89. The Atlas Centre and its staff (although components of the central IT team were located at various times in other buildings on site) achieved much: from superb 24/7 operations of a major computer centre over many years and with many different kinds of equipment (now 1100 servers and 360000 users); provision of services in administration and management for the organisation (and others) front-ended by office systems, directories etc; provision of systems for other scientific departments, the library and others; development of websites and behind it all a powerful R&D capability from graphics, visualisation and virtual reality through databases, information retrieval and office systems to formal software engineering, knowledge engineering, trust/privacy/security and onwards via advanced user interfaces and management to applied mathematics and computational optimisation. From this team came the original UK e-Science concept. Along the way companies were spun out and contibutions made to standards notably OSI networking, GKS, SVG and SMIL.

# Albert Benveniste wins the "France Telecom" 2008 Award

Albert Benveniste, a world-renowned expert in the fields of automatic control, computer science and telecommunications, has been granted the "France Telecom" 2008 award by the French Academy of Sciences for the exceptional quality of his research work. Director of Research at INRIA Rennes - Bretagne Atlantique, Albert Benveniste is responsible for the embedded systems sector at the institute's scientific department. He is also scientific director of the INRIA - Alcatel-Lucent Bell Labs joint research laboratory.

The originality of Albert Benveniste's profile lies in the extraordinary breadth of his expertise and the foresightedness and depth of his contributions. In the 1980s, for example, he was the first person to recognise the importance of jointly managing algorithmic and software aspects for the design of telecommunication systems and embedded systems. His ongoing, innovative activities at both the national and international levels have contributed significantly to cross-fertilisation between the automatic control, signal processing and computer science communities. This unique profile puts him in a perfect position to tackle the fields of telecommunication networks and services management as well as Web services, subjects in which difficult software and algorithmic aspects are closely linked and which are at the heart of today's challenges.

ERCIM
European Research Consortium for Informatics and Mathematics

**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.**

AARIT
Austrian Association for Research in IT
c/o Österreichische Computer Gesellschaft
Wollzeile 1-3, A-1010 Wien, Austria
http://www.aarit.at/

CNR
Consiglio Nazionale delle Ricerche, ISTI-CNR
Area della Ricerca CNR di Pisa,
Via G. Moruzzi 1, 56124 Pisa, Italy
http://www.isti.cnr.it/

Crcim
Czech Research Consortium
for Informatics and Mathematics
FI MU, Botanicka 68a, CZ-602 00 Brno, Czech Republic
http://www.utia.cas.cz/CRCIM/home.html

CWI
Centrum Wiskunde & Informatica
Centrum Wiskunde & Informatica
Science Park 123, NL-1098 XG Amsterdam,
The Netherlands
http://www.cwi.nl/

danaim
Danish Research Association for Informatics and Mathematics
c/o Aalborg University,
Selma Lagerlöfs Vej 300, 9220 Aalborg East, Denmark
http://www.danaim.dk/

Fonds National de la Recherche Luxembourg
Fonds National de la Recherche
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
http://www.fnr.lu/

FWO
Egmontstraat 5
B-1000 Brussels, Belgium
http://www.fwo.be/

FNRS
rue d'Egmont 5
B-1000 Brussels, Belgium
http://www.fnrs.be/

FORTH
Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
http://www.ics.forth.gr/

Fraunhofer Gesellschaft
Fraunhofer ICT Group
Friedrichstr. 60
10117 Berlin, Germany
http://www.iuk.fraunhofer.de/

INRIA
Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
http://www.inria.fr/

IUA
Irish Universities Association
Irish Universities Association
Cumann Ollscoileanna Éireann
Irish Universities Association
c/o School of Computing, Dublin City University
Glasnevin, Dublin 9, Ireland
http://ercim.computing.dcu.ie/

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and
Electrical Engineering, N 7491 Trondheim, Norway
http://www.ntnu.no/

PEG
PORTUGUESE ERCIM GROUPING
Portuguese ERCIM Grouping
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, nº 378,
4200-465 Porto, Portugal

PLERCIM
Universitas Varsoviensis
Universitas Wratislaviensis
Polish Research Consortium for Informatics and Mathematics
Wydział Matematyki, Informatyki i Mechaniki,
Uniwersytetu Warszawskiego, ul. Banacha 2, 02-097 Warszawa, Poland
http://www.plercim.pl/

Science & Technology Facilities Council
Science and Technology Facilities Council,
Rutherford Appleton Laboratory
Harwell Science and Innovation Campus
Chilton, Didcot, Oxfordshire OX11 0QX, United Kingdom
http://www.scitech.ac.uk/

SpaRCIM
Spanish Research Consortium for Informatics
and Mathematics c/o Esperanza Marcos, Rey Juan Carlos University,
C/ Tulipan s/n, 28933-Móstoles, Madrid, Spain,
http://kybele.escet.urjc.es/SpaRCIM/

SICS
SWEDISH INSTITUTE OF COMPUTER SCIENCE
Swedish Institute of Computer Science
Box 1263,
SE-164 29 Kista, Sweden
http://www.sics.se/

SARIT
Swiss Association for Research in Information Technology
c/o Professor Daniel Thalmann, EPFL-VRlab,
CH-1015 Lausanne, Switzerland
http://www.sarit.ch/

MTA SZTAKI
Magyar Tudományos Akadémia
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
http://www.sztaki.hu/

VTT
Technical Research Centre of Finland
PO Box 1000
FIN-02044 VTT, Finland
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 edition*** ☐ ***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 Web site at http://ercim-news.ercim.org/**