


# ERCIM



# NEWS

European Research Consortium  
for Informatics and Mathematics  
[www.ercim.org](http://www.ercim.org)

Special theme:

# Future Internet Technology

## Also in this issue:

*Keynote*  
by Viviane Reding

*Joint ERCIM Actions:*  
Twenty Years of ERCIM:  
History and Outlook

*R&D and Technology Transfer:*  
Fighting Lymph Cancer  
with Mathematics

*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](mailto:contact@ercim.org)  
Director: Jérôme Chailloux  
ISSN 0926-4981*

**Editorial Board:**

*Central editor:*

*Peter Kunz, ERCIM office ([peter.kunz@ercim.org](mailto:peter.kunz@ercim.org))*

*Local Editors:*

*Austria: Erwin Schoitsch, ([erwin.schoitsch@arcs.ac.at](mailto:erwin.schoitsch@arcs.ac.at))*

*Belgium: Benoît Michel ([benoit.michel@uclouvain.be](mailto:benoit.michel@uclouvain.be))*

*Denmark: Jens Bennedsen ([jbb@it-vest.dk](mailto:jbb@it-vest.dk))*

*Czech Republic: Michal Haindl ([haindl@utia.cas.cz](mailto:haindl@utia.cas.cz))*

*France: Bernard Hidoine ([bernard.hidoine@inria.fr](mailto:bernard.hidoine@inria.fr))*

*Germany: Michael Krapp ([michael.krapp@scai.fraunhofer.de](mailto:michael.krapp@scai.fraunhofer.de))*

*Greece: Eleni Orphanoudakis ([eleni@ics.forth.gr](mailto:eleni@ics.forth.gr))*

*Hungary: Erzsébet Csuha-Varjú ([csuhaj@sztaki.hu](mailto:csuhaj@sztaki.hu))*

*Ireland: Ray Walsh ([ray@computing.dcu.ie](mailto:ray@computing.dcu.ie))*

*Italy: Carol Peters ([carol.peters@jsti.cnr.it](mailto:carol.peters@jsti.cnr.it))*

*Luxembourg: Patrik Hitzelberger ([hitzelbe@lippmann.lu](mailto:hitzelbe@lippmann.lu))*

*Norway: Truls Gjestland ([truls.gjestland@ime.ntnu.no](mailto:truls.gjestland@ime.ntnu.no))*

*Poland: Hung Son Nguyen ([son@mimuw.edu.pl](mailto:son@mimuw.edu.pl))*

*Portugal: Paulo Ferreira ([paulo.ferreira@inesc-id.pt](mailto:paulo.ferreira@inesc-id.pt))*

*Spain: Christophe Joubert ([joubert@dsic.upv.es](mailto:joubert@dsic.upv.es))*

*Sweden: Kersti Hedman ([kersti@sics.se](mailto:kersti@sics.se))*

*Switzerland: Harry Rudin ([hrudin@smile.ch](mailto:hrudin@smile.ch))*

*The Netherlands: Annette Kik ([Annette.Kik@cw.nl](mailto:Annette.Kik@cw.nl))*

*United Kingdom: Martin Prime ([M.J.Prime@stfc.ac.uk](mailto:M.J.Prime@stfc.ac.uk))*

*W3C: Marie-Claire Fogue ([mcf@w3.org](mailto:mcf@w3.org))*

**Contributions**

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

**Copyright Notice**

*All authors, as identified in each article, retain copyright of their work*

**Advertising**

*For current advertising rates and conditions, see*

*<http://ercim-news.ercim.org/> or contact [peter.kunz@ercim.org](mailto: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](mailto:en-subscriptions@ercim.org)*

*or by filling out the form at the ERCIM News website:*

*<http://ercim-news.ercim.org/>*

**Next issue:**

*July 2009, Special theme: Mathematics for Finance and Economics*



*Viviane Reding, Member of the European Commission  
Information Society & Media*

The Internet is undoubtedly permeating and transforming all aspects of our economies and societies. It is a remarkable catalyst for creativity, collaboration and innovation and more broadly, for the development of our economies and societies. A few examples: in 1998, Google indexed 26 million web-pages, today it indexes 1 trillion; Within only five years, Facebook and MySpace have attracted each more than 100 million users world wide; user-generated content such as YouTube produced more than 73 billion streams in 2008; with around four billion mobile users world wide, the Internet is becoming more and more mobile and capable to support a range of new applications and services that were not foreseen in its original design.

Contrary to the common-sense “intuitive linear” view of progress, Ray Kurzweil predicts accelerating returns: at today’s rate, we will experience as much progress in the next 100 years as we saw in the past 20,000 years! This revolution will only become more evident as the Internet becomes more pervasive, available anywhere and at anytime. By 2011, we expect that about 3 billion hosts will be connected to the Internet, (a more than 500% increase compared to 2008) and that by 2012 annual global Internet traffic will reach ½ zettabyte, that is, 250,000 times more than in 2003.

In the coming years we expect an exponential increase of information, the further development of social networks, the accelerated growth of online video traffic, and the emergence of the Internet of things, which will progressively cause the online and real worlds to become inter-linked. Mobility and extended nomadicity, the growth of security and privacy threats, the diversity and sheer number of new applications, services and business models supported by the Internet mean that the existing Internet architecture may soon be placed under strain.

This raises the crucial question as to whether the today’s Internet architecture needs to be redesigned and, in particular, how should it evolve in order to meet such challenges as: availability, heterogeneity, scalability, mobility, manageability, security, trust, openness and neutrality; while at the same

time guaranteeing that business and governance models are sustainable and supporting spontaneous and emerging behaviour, user creativity and unanticipated new usages.

These and other dimensions of the future Internet need to be carefully explored. The research community has a key role to play in this exploration to shape a future of growth and prosperity in Europe.

With the 7th research framework programme, the European Union has launched more than 90 Information and Communication Technologies projects addressing several challenging dimensions of the future Internet. These projects are rallying over 500 organisations across Europe - and have an overall budget of more than 400 million Euros. This represents one of the most significant investments ever made in Internet-related research which will propel Europe into a leading worldwide position.

However, the sheer size of the challenges ahead requires the coordination of efforts of everyone involved in developing Europe’s future Internet. The creation of the Future Internet Assembly (FIA) which gathers all major European actors has been a first decisive step.

I am pleased to note the FIA’s all-embracing approach that combines a multi-stakeholder industrial strategy, visionary scenarios and multidisciplinary research underpinning a new generation of Internet technologies while supporting shared European values such as social acquis, culture and inclusion.

The European Commission is also supporting the future Internet through several policy-related initiatives. Our recent Communication on future networks and the Internet highlights the need and opportunities presented by launching a public-private partnership at European level that would consolidate the currently fragmented research efforts in this field.

All these European actions support our ambition to strengthen Europe’s footprint on the Internet of tomorrow to shape its technological and socio-economic development. Only by working together, can we assure that Europe gains the competitive advantage it needs to succeed in the 21st century economy.

This issue of ERCIM News provides a comprehensive overview of the numerous activities, projects and initiatives addressing Future Internet research all around Europe. I welcome the remarkable achievements and ambitious goals illustrated in this report which prove that European academic and industrial stakeholders are well aware and well prepared to address the opportunities and challenges ahead.

*Viviane Reding*

## 2 Editorial Information

### KEYNOTE

- 3 by Viviane Reding, Member of the European Commission, Information Society & Media.

### JOINT ERCIM ACTIONS

- 6 **Twenty Years of ERCIM: History and Outlook**  
by Keith Jeffery
- 6 **Twenty Years of ERCIM: Testimonials**
- 10 **Digital Libraries: The pioneering Role of ERCIM in establishing this Research Field in Europe**  
by Costantino Thanos
- 11 **Formal Methods for Industrial Critical Systems: ERCIM's oldest active Working Group**  
by Alessandro Fantechi
- 12 **Twenty Years of ERCIM News**
- 13 **ERCIM Working Group 'User Interfaces for All' European Coordination Actions: Beyond-The-Horizon & InterLink**  
by Constantine Stephanidis
- 14 **Eye to the Future – EchoGRID Roadmap: A Shared Vision for EU-China Cooperation on Next-Generation Grid Research**  
by Stephanie Parker
- 14 **EuroIndia International Events**  
by Ashok Kar, Sudhanshu Rai and Florence Pesce

## SPECIAL THEME

### Introduction to the Special Theme

- 16 **Future Internet Technology**  
by François Baccelli and Jon Crowcroft

### Setting The Context

- 19 *Invited article*  
**Starting the Debate: Agreeing on Disagreements**  
by Dirk Trossen

20 **The Internet Engineering Task Force and the Future of the Internet**

by Emmanuel Baccelli, Thomas H. Clausen and Philippe Jacquet

### Architecture

- 21 *Invited article*  
**Network Virtualization - An Enabler for Overcoming Ossification**  
by Anja Feldmann, Mario Kind, Olaf Maennel, Gregor Schaffrath and Christoph Werle

23 **Y-COMM: A New Architecture for Mobile Heterogeneous Communications**

by Glenford Mapp

### Modelling, Measurement and Management

- 24 *Invited article*  
**Epidemic Information Dissemination**  
by Laurent Massoulié

25 **MyNetScope: A Platform for Network Traffic Visualization and Analysis**

by Pavel Minařík

26 **Unified Access to Internet Measurement Data**  
by Felix Strohmeier, Martin Nilsson and Demetres Antoniadis

28 **Privacy-Aware Network Monitoring**

by Peter Dorfinger, Carsten Schmoll and Felix Strohmeier

29 **Breaking the Weakest Link: Becoming a Trusted Authority on the Internet**

by Marc Stevens

### Correctness and Robustness

- 30 **ResiliNets: Resilient and Survivable Networks**  
by David Hutchison and James P.G. Sterbenz

32 **Standardized Testware for Internet-Based Telecommunication Services**

by Bostjan Pintar, Axel Rennoch, Peter Schmitting and Stephan Schulz

33 **Network Description Tools and Standards**

by Freek Dijkstra, Jeroen van der Ham and Ronald van der Pol

### Novel Applications

#### 34 *Invited article*

**Delay-Tolerant Bulk Internet Transfers**  
by Nikolaos Laoutaris and Pablo Rodriguez

**36 Resource Management for IPTV Distribution**  
by Henrik Abrahamsson and Per Kreuger

**37 VISTO: Visual Storyboard for Web Video Browsing, Searching and Indexing**  
by Marco Pellegrini

**38 The New Role of Humans in the Future Internet**  
by Daniel Schall and Schahram Dustdar

### In the Home

**40 Orchestrating your Surroundings**  
by Pau Giner, Carlos Cetina, Joan Fons and Vicente Pelechano

**41 An Autonomic Home Networking Infrastructure**  
by Thomas Luckenbach, Mario Schuster and Marc-Oliver Pahl

**42 Load-Balancing Energy Usage of Household Appliances**  
by Lennart E. Fahlén

### On the Road

**43 The Internet of Vehicles or the Second Generation of Telematic Services**  
by Markus Miche and Thomas Michael Bohnert

**45 Web-Enabled Tracking Operations in Distributed Supply Chains**  
by Zsolt Kemény and Elisabeth Ilie-Zudor

### Awareness, Identity and Society

**46 Knowledge-Based Collaboration Patterns in Future Internet Enterprise Systems**  
by Gregoris Mentzas and Keith Popplewell

**48 Experiences from the Public Safety Arena for the Future Internet**  
by Socrates Varakliotis, Peter Kirstein and Steve Hailes

**49 With Joint Forces - Establishing Situation Awareness for the Future Internet**  
by Tanja Zseby and Thomas Hirsch

**51 Social Networking for a Pervasive Future Internet: the SOCIALNETS Project**  
by Stuart M. Allen, Marco Conti, Andrea Passarella and Roger M. Whitaker

**52 Owner-Centric Networking: A New Architecture for a Pollution-Free Internet**  
by Claude Castelluccia and Mohamed Ali Kaafar

**53 Semantic Web for e-Commerce**  
by Bernd Gruber

### Wireless, Embedded

**55 Management of Future Mesh-Based Radio Access Networks**  
by Vangelis Angelakis, Vasilios Siris and Apostolos Tragantitis

**56 Testing Mobile Data Applications on Smartphones with SymPA**  
by Almudena Díaz and Pedro Merino

## R&D AND TECHNOLOGY TRANSFER

**58 Fighting Lymph Cancer with Mathematics**  
by Annette Kik and Gunnar Klau

**59 Measuring Digital Library Usage Using Network Traffic Analysis**  
by Jiří Šmerda and Radka Findeisová

**60 Multi-Faceted Rating of Product Reviews**  
by Stefano Baccianella, Andrea Esuli and Fabrizio Sebastiani

**61 More on Mobile Telephones and Our Health**  
by Harry Rudin

**62 The Open Wall: A Software-Intensive Art Installation**  
by Letizia Jaccheri

**64 Simulating Invisible Dangers in Emergencies: The SimRad-Projects**  
by Gerhard Chroust, Karin Rainer and Markus Roth

**65 DL.org: A Coordination Action on Digital Library Interoperability, Best Practices and Modelling Foundations**  
by Donatella Castelli and Stephanie Parker

## EVENTS

**66 FLReNet: Shaping the Future of a Multilingual Digital Europe**  
by Nicoletta Calzolari

**67 1st Workshop on Formal Methods for Service-Oriented Architecture and Internet of the Future**

**67 Announcements**

**71 In Brief**



Photo: Stigfried Münch

*The 'birth' of ERCIM: From left - Alain Bensoussan, INRIA, Friedrich Winkelhage, GMD (on behalf of Gerhard Seegmüller) and Cor Baayen, CWI, sign the joint agreement on 13 April 1989.*



Photo: INRIA

*The ERCIM 'founding fathers' Gerhard Seegmüller (GMD), Alain Bensoussan (INRIA) and Cor Baayen (CWI) in 1989.*



Photo: CWI

*The first 'ERCIM fellows' Michal Haindl (left) and Eric Rutten (right) during an ERCIM workshop at CWI in Amsterdam in November 1990 (in the center Chris Greenough of RAL, now STFC). More than 200 fellows have been hosted in ERCIM institutes through the Fellowship Programme since its inception in 1990.*

## Twenty Years of ERCIM: History and Outlook



*Keith Jeffery,  
President of ERCIM.*

*ERCIM is celebrating twenty years of cooperation for excellence in research. Over the last twenty years, ERCIM has developed from the initial three-member consortium to an open and inclusive - but managed - European organization of currently twenty members.*

ERCIM is based on a model of one member per country, each acting as a node linking academia and industry in that country and – via the other ERCIM members – to other countries. ERCIM can be envisaged as a high-speed linking network connecting and supporting researchers across Europe. ERCIM is open to anyone wishing to participate in the thematic working groups and the related workshops and R&D projects, or wishing to publish in ERCIM News, all via the relevant national member or the Office.

Based upon this robust structure, ERCIM has helped develop and support European-centric multinational communities of high quality in particular subject domains within ICT and applied mathematics. These communities have achieved significant, high-quality, cooperative research across academia and industry from 'blue sky' research to highly applied development. To take just a few examples in which ERCIM has played a major role: digital libraries (DELOS, DILIGENT, D4Science, DRIVER); Grids (CoreGRID, GRIDCOMP, TrustGRID, ECHOGRID); WWW (projects leading to SMIL standard, SVG standard and early work on database/Web interfaces); trust Security and privacy (TRUSTCOM, i-TRUST, GRIDTRUST, PRIME); heterogeneous distributed database integration; software engineering including formal methods across embedded and wide-scale platforms; and user interfaces (UI4ALL, SESAME). This has extended into many application areas, examples being: biomedical (ACGT, VPH); environmental (DECAIR, THETIS, SIMES, CRUCID, TELEMAT, ANFAS); educational, cultural and multimedia (AQUARELLE, VITALAS, MUSCLE, PALETTE).

ERCIM has put in place an ambitious postdoctoral fellowship programme hosted across ERCIM institutes, to increase European cohesion and the supply of well-trained and experienced researchers. The Cor Baayen award for a promising young is given based on open competition Europe-wide and thus encourages young researchers.

ERCIM focuses the experience and ideas from a wide European constituency and, informed also by international inter-

actions outside Europe, provides advice on strategic directions in ICT and applied mathematics to the European Commission (BEYOND-THE-HORIZON, INTERLINK) and other organizations. ERCIM has undertaken consultancy contracts for the EC, and has acted as the EU representative in workshops and discussions with the USA NSF (National Science Foundation). Past activities have involved the Mediterranean region, with more recent activities focusing on India and China. Senior ERCIM researchers have acted as experts for the EC on strategic issues (eg GRIDs) and as reviewers.

ERCIM News – which is very widely read (in print and online) and which attracts influential personalities to write the foreword of each issue - has become the ‘must read’ publication in ICT and applied mathematics R&D in Europe (and beyond). The ERCIM Web site has a high page ranking and is much visited, giving ERCIM a high publicity profile.

To ensure the relevance and quality of our work, ERCIM has set up an external advisory committee of independent highly respected members of the international community to advise the President and Board of Directors. ERCIM has set up strategic partnerships with ETSI (European Telecommunications Standards Institute) and EMS (European Mathematical Society), with early tangible results (the Infinity Initiative seminar series and a special issue of ERCIM News respectively) and evolving plans evident in each case.

ERCIM took over from INRIA the responsibility as European host of the W3C (World Wide Web Consortium), and is contributing strongly to the development of the Web in both technology leading to standards and the increasing utilization and take-up in Europe through awareness, education and support.

ERCIM has in recent years emphasised innovation, and knowledge and technology transfer, building upon the more than 200 spin-out companies from ERCIM members and associated products and patents. A strategic seminar programme engages with industry on topics in which ERCIM can offer research and development expertise. A new publication, ‘ERCIM Innovation’, showcasing ERCIM R&D offerings but also containing articles on related issues such as venture capital and market analysis, was launched at the EC ICT2008 Conference. A network of innovation and knowledge transfer experts among ERCIM members, overlaid on the network of researchers focused through ERCIM, assists in the Europe-wide take-up of research from any one (or more) members for wealth creation and improvement in the quality of life. Furthermore, our work with W3C and ETSI allows early standardization of R&D, leading to a larger stable market for products and services.

While achievements mentioned above provide the bare facts, behind these achievements are the people who make ERCIM what it is. I have the honour of following in the footsteps of a sequence of distinguished presidents, each of whom brought

*“It gives me great pleasure and satisfaction to observe that ERCIM succeeds so well in realising the objectives of its founders: creating coherence in Europe in research in mathematics and information science through personal acquaintance and co-operation of researchers from different nations, and at the same time, using these contacts to stimulate excellence in research and innovation in the participating national organisations.”*

*More than ever this co-operative excellence in the key areas covered by ERCIM is necessary to help counter and redress the problems of our society. I firmly believe that the joint research activities of ERCIM will prove to be essential.”*

*Cor Baayen, First ERCIM President (1989 - 1994) and ERCIM Président d’honneur*



Photo: CWI



Photo: INRIA

*“It is a real pleasure to write a few words about ERCIM. Cor Baayen (CWI), Gerhard Seegmüller (GMD) and I created ERCIM twenty years ago, but initiated it after the Single European Act of 1987. In those days some kind of Euphoria was spreading in Europe, with the impression that a significant step towards European Integration had*

*been attained. Our simple question was this: How can we contribute to this general trend?*

*ERCIM has been the first EEIG for research and in some way anticipated the Research European networks, created by the successive framework programs. We also wanted a consortium including Eastern and Central Europe, at a time when this was still a remote dream.*

*To see today the successful path of ERCIM is like seeing your child becoming a successful adult, with a solid position and a rich future. It is a real pleasure, a real pride.*

*My best wishes to ERCIM for 20 new successful years, and my gratitude to our successors who have done a terrific job.”*

*Alain Bensoussan, President of INRIA from 1984 to 1996*



*“It has been a delight to serve as ERCIM’s first Advisory Committee (AC) Chair. I found ERCIM’s unique blend of research vision and innovation leadership to be a major driving force to shape Europe’s future and the world at large. Further, I, along with all AC members found ERCIM to embody the ideals of a well-established organization. The energy of its leaders, executives and staff is outstanding, and always willing to go the extra mile to incorporate the feedback they receive from the AC.”*

*Mazin Yousif, PhD, Avirtec, Inc., CEO and Ex-Intel Executive ERCIM Advisory Committee Chair*



Photo: CWI

*ERCIM Directors give a toast on Rutherford Appleton Laboratory (RAL) joining ERCIM as fourth ERCIM member in November 1990. From left: Alain Bensoussan (INRIA), Gerhard Seegmuller (GMD), Paul Williams (RAL), Cor Baayen (CWI).*




Photo: Fred Schuurber, VU/CWI

*ERCIM celebrated its 10th anniversary in Amsterdam in November 1999. ERCIM president Gerard van Oortmerssen (center) and Stelios Orphanoudakis who succeeded him as president in 2004, are blowing the candles. Sadly, due to illness, Stelios' presidency was all too short.*

their own energy and style to the advancement of ERCIM. We now have a clear management structure. The Board of Directors considers strategic issues and represents the top management of the member organizations. It has Strategy Task Groups on Structure, External Relations and Innovation, each led by a vice-president. The Executive Committee manages the operational aspects of ERCIM and has Task Groups covering Finance, Projects, Human Resources and Public Relations. The Executive Committee coordinates three important aspects of ERCIM: the Working Groups, which run workshops, build communities and have research projects; the Editorial Board, which produces ERCIM News; and the ERCIM Office, which supports all areas of ERCIM activity including the management of projects.

The past twenty years have seen amazing advances. We have moved from kilobit to gigabit communication bandwidths.

*Evolution of ERCIM. The graph shows the increase in ERCIM Membership, Working Groups (WG), European contracts managed by ERCIM and Fellows hosted at ERCIM institutes through the 'Alain Bensoussan' Fellowship Programme.*



*Dennis Tsichritzis, President of ERCIM (1994-1998) and the Board of Directors welcome the President of the Czech Research Consortium for Informatics and Mathematics, Jiri Wiedermann, during the ERCIM meetings in Zurich in May 1996.*



We have seen the widespread adoption of personal devices from laptops to mobile phones connected anywhere. Sensor networks have appeared and data volumes have increased dramatically. Hardware costs have decreased and software costs increased. Systems and software engineering techniques have developed relatively little. The World Wide Web provided a stunningly simple but brilliant solution to sharing information. Linked with database and information retrieval technology for managing information, and with graphics for presenting information, it has revolutionised ICT. The early promise of AI (Artificial Intelligence) has not been realised, but robotics has developed considerably and some aspects of knowledge engineering (eg domain ontologies) have had demonstrable successes. Above all, the widespread use of ICT (and the underlying applied mathematics) in all contexts has been startling. The billions of processors embedded in everyday devices, the massive use of ICT for social interactions and the ever-increasing use of ICT for professional activities has amazed everyone.

“Prediction is very difficult, especially if it’s about the future.” (Niels Bohr). However, at such a milestone it is more or less obligatory. It is clear that despite dire predictions, the limits of current processor and memory technology have not yet been reached and are being pushed further. On the other hand, newer hardware technologies are being developed in parallel with more speculative advances in quantum and biological computing. A connected world of ever-increasing numbers of sensors will also take shape. There are implications for the balance of media types: audio and video will become more prevalent.

The future of software is much less clear. The move towards self-describing services that can be combined in a ‘plug and play’ mode (either by the software engineer or automatically) seems inevitable and is driven by factors of cost, reliability, security, durability and flexibility. Such a move should improve reuse of components and interoperability of application systems, with clear advantages for management.

The merging of communications, location (spatial and temporal), media and processing will parallel the development of a seamless personal lifestyle across professional and social contexts, but we must be vigilant to the dangers to security and privacy. Intelligent environments in the home, office, factory, control room, medical centre, vehicles and so on will support this. Information supplied on demand by the user will evolve towards information pushed and filtered to the user – for whatever context and role is appropriate. User interfaces will move through haptic to ‘brain-linked’.

As it has over the last two decades, I am confident that ERCIM will remain at the forefront of developments into the future. Plans have been put in place for this to happen effectively and efficiently, and the people – young dynamic excellent researchers – are positioned to take over from the current senior generation. ERCIM has a great future.

*Keith Jeffery, ERCIM President*



Photo: ICTPage

*“I cherish my memories of my involvement in ERCIM, from 1991 to 2003. Especially my period as president, during which we celebrated the 10th anniversary in Amsterdam, and accomplished the transfer of the European W3C Office to ERCIM, was exciting and rewarding. The consortium strengthened its ties with the European Commission, the European Science Foundation and the US National Science Foundation and promoted its role as a distributed think tank in information and communications technology and applied mathematics. ICT is a driver of innovation, and European cooperation is crucial for real progress. ERCIM is therefore a very valuable network. After I left CWI and thereby the ERCIM community, I still followed the developments with much interest, and I look forward to each new issue of ERCIM News. Usually I am pleasantly surprised by the ability of ERCIM to address novel topics which are of interest to researchers as well as to industry and society.”*

*Gerard van Oortmerssen,  
ERCIM President from May 1998 to December 2003*



*“My relationship with ERCIM is a long story. It started even before ERCIM’s birth and thus more than twenty years ago, when I was in charge of international relations for INRIA. At that time, three European research institutes in mathematics and computer science, INRIA, CWI and GMD (later merged with Fraunhofer Gesellschaft), used to cooperate in the framework of bilateral agreements. Once, instead of holding their annual bilateral meetings, their Directors decided to have a joint meeting at GMD headquarters in Sankt Augustin. After a morning of talks, they asked both my German and Dutch counterparts as well as me to provide them with the project of a trilateral agreement after lunch time. It was a sunny day. We sat around a table, forgot about our lunch, and were able to present a first draft two hours later. As far as I know this was the first step towards a European consortium in informatics and mathematics.*

*When I came back to INRIA, after a few years spent at the French Ministry of Research, ERCIM had been setup. It was a young and promising body of which I had the pleasure of being appointed as Manager. The team was small but very efficient. This was the first EEIG in the field of research and we felt participating in building something new. We were strongly supported by the Board of Directors and a proactive Executive Committee. We had the chance to win several European projects and we contributed to increase ERCIM’s visibility on the international scene. This was a challenging and most fruitful experience and I keep a very vivid memory of these two years within ERCIM.”*

*Jean-Michel Chasseriaux,  
ERCIM Manager from 1994 to 1996*

## Digital Libraries: The pioneering Role of ERCIM in establishing this Research Field in Europe

by Costantino Thanos

*ERCIM has played a pioneering role in establishing the Digital Library (DL) field as a distinct area of research in Europe. Instrumental have been, first, the activities of the ERCIM Working Group on Digital Libraries, set up in 1995, and later those of the DELOS Network of Excellence on Digital Libraries.*

The main objectives of the ERCIM WG on Digital Libraries were to stimulate interest and to encourage collaboration between the ERCIM research teams beginning to be active in this field. A first result was the development of the ERCIM Technical Reference Digital Library (ETRD) which provided an experimental DL service for ERCIM scientists.

An initiative full of future fruitful consequences of the ERCIM WG was the creation of the DELOS Working Group on Digital Libraries funded by the EU ESPRIT Long Term Research Programme under the 4th Framework Programme (1997 – 1999). The objectives of the DELOS WG were to arouse awareness of priority DL research issues, to encourage collaboration between European research teams, and to establish links with on-going international projects in the DL field. An important achievement was the establishment of a formal collaboration with the US National Science Foundation and the creation of five joint EU-US collaborative Working Groups. These WGs explored DL-related technical, social and economic issues, and published a set of recommendations with respect to DL interoperability, metadata, IPRs and economics, global resource discovery and multilingual information access in a special issue of the International Journal of Digital Libraries.

An important spin-off of the DELOS WG activities was the launching of the European Conference on Digital Libraries – ECDL. The annual ECDL conference rapidly became the main forum for the European DL community to present and discuss their research ideas and results. The 13th edition of the conference will be held this year in Corfu, Greece, 27 September – 02 October 2009.

In 2000, the DELOS WG gave way to the DELOS Thematic Network of Excellence on Digital Libraries, funded by the EU under FP5 (2000 – 2003). The activities of the Network included the organization of workshops, summer schools, important evaluation initiatives (CLEF, INEX), and technology transfer actions. In addition seven DELOS – NSF joint WGs were created: DL Infrastructures; Digital Imaging for Cultural and Historical Materials; Digital Archiving and Preservation; Personalization and Recommender Systems;

Actors in DLs; Spoken-word Digital Audio Collections; and Emerging Language Technologies and Rediscovery of the Past.

In a Roadmap Report the DELOS research community defined the DL Grand Vision: “Digital Libraries should enable any citizen to access all human knowledge any time and anywhere, in a friendly, multimodal, efficient and effective way, by overcoming barriers of distance, language, and culture and by using multiple Internet-connected devices”.

The last phase of the DELOS evolution was its transformation into the DELOS Network of Excellence, under FP6 (2004 – 2007). The mission was to integrate and coordinate the on-going research activities of the major European research teams in the field of Digital Libraries. The main



*Costantino Thanos presents the ‘DELOS award’ for the best paper written by a young researcher at the 10th edition of the ECDL conference in Alicante 2006. ECDL - the European Conference on Digital Libraries was a ‘spin-off’ of the ERCIM DELOS Working Group. This year, the 13th edition of the conference will be in Corfu, Greece, from 27 September to 2 October 2009.*

achievement was the definition of the “DELOS DL Reference Model”, a formal and conceptual framework describing the characteristics of Digital Library Management Systems and the development of a Reference Architecture for such systems.

Under FP4, FP5 and FP6, DELOS has significantly influenced research, development, and educational activities in the DL field in Europe:

- At the research level, a newly funded EU project, DL.org, aims to further develop and consolidate the DELOS DL Reference Model with the ultimate goal of making it an internationally accepted standard.
- At the development level, DELOS has inspired many EU projects aiming at developing scientific repositories and Digital Library infrastructures (Scholnet, Cyclades, Diligent, D4Science).
- At the educational level, the DELOS DL Reference Model is taught in several European DL university courses.

In addition, significant expertise created under DELOS has been canalized into "Europeana", the major effort of the European Commission towards the building of the European Digital Library.

The main merits of DELOS, with the help and support of ERCIM, are to have significantly contributed to the creation of a European DL research community, to have contributed to create important durable scientific events and infrastructures (ECDL, CLEF) and to have anticipated, in the form of either prototypes or roadmap reports, many actions of the European Commission in the field of Digital Libraries.

## Formal Methods for Industrial Critical Systems: ERCIM's oldest active Working Group

by Alessandro Fantechi

Following an initial successful workshop bringing together ERCIM members interested in formal verification, held in Pisa in December 1992, Stefania Gnesi and Diego Latella, CNR, Pisa, proposed the creation of an ERCIM Working Group dedicated to Formal Methods for Industrial Critical Systems (FMICS). Although at that time, model checking was in its early days, the early ERCIM FMICS community was already aware of the great potential of formal verification techniques.

Since then, the WG, chaired in succession by Latella (ISTI-CNR), Hubert Garavel (INRIA), Gnesi (ISTI) and Pedro Merino (SpaRCIM), has kept pace with the development of formal verification techniques - and model checking in particular. The series of annual workshops, begun in 1996 and sponsored by the WG, have promoted an ongoing scientific discussion focussed on identifying the most efficient verification techniques, with a keen eye on their industrial applicability. Most of the members of the FMICS community have strong links with industry and have thus contributed to



ERCIM President Gerard van Oortmerssen presents the Working Group award 2002 to Stefania Gnesi, coordinator of the Working Group on Formal Methods for Industrial Critical Systems.

*"In 1990, the three initial ERCIM members decided to expand the membership to have a broader representation across Europe. Rutherford Appleton Laboratory (RAL) was the first new member and I joined an informal Executive Committee that spent quite a bit of time on the road assessing the various candidates for membership.*

*In the next two years, the membership rose from four to twelve and there was a need to introduce a more formal Executive Committee reporting to the ERCIM Directors. I chaired the Committee from 1992 to 1995. It was an exciting time. Some of our achievements were:*

- *the individual members did get to know each other better through the workshops that were organised. Joint participation took place in projects that probably would not have happened previously*
- *ERCIM made a major impact on the mobility of researchers across Europe through the Fellowship scheme that was launched and the Cor Baayen Award*
- *ERCIM News was launched and it is pleasing to see how it has developed into a major European scientific publication over the years*
- *specific research areas were influenced significantly by the ERCIM members. Some that come immediately to mind are electronic libraries, parallel processing, databases, and, of course, the World Wide Web.*

*I believe ERCIM played a major part in the establishment of a strong European Web presence early on. INRIA became the European Host when CERN withdrew and the ERCIM members seconded staff to INRIA to help in the initial launch of the Office. The W3C European Offices were established at ERCIM member sites and the raising of awareness in Europe was facilitated by the fact that the Offices knew each other through ERCIM.*

*By the time I left the Committee in 1995, ERCIM was a significantly different organisation from the one that RAL had joined a few years earlier, larger, more confident, more influential. It is a pleasure to see how it has advanced over the years."*

*Bob Hopgood, ERCIM Executive Committee chair (1992 - 1995)*



Georges Nissen, ERCIM Executive Committee member (1989 - 2001) remembers how everything started:

*"In the late eighties, INRIA researchers had many point-to-point scientific collaborations with other European researchers in Universities or Research Institutes. The links with*

*CWI were particularly active, due to a common scientific culture, namely the coexistence of two principal fields: Computer Science (improving the tool) and Scientific Computation (improving the principal applications of the tool). These fields are often separate, and this coexistence is obviously a potential source of synergy in both organisations.*

*These excellent 'grassroot' relations oozed upwards in the Ministerial world of both countries, and one of their yearly meetings came to the conclusion of 'twinning' both organisations. Somewhat anecdotically, our Dutch partner came back with this information a few months later, whereas on our side, this information got lost. Anyway, we could start working together more formally (launching workshops, exchanging scientists, etc.). The next step, a year and a half later, was to bring in GMD which had quite similar scientific interests in a somewhat different structure, and with which INRIA had political links since quite a few years. The gathering process was launched"*

the slow, but constant introduction of formal methods in the development cycle of industrial critical systems witnessed in the last decade.

The WG has also addressed other readily applicable verification techniques, such as static analysis by abstract interpretation. Similarly, the whole formal development life cycle has been addressed, for example in the 2008 FMICS workshop, where considerable attention was paid to the recent diffusion of Model Driven Development in industry.

ERCIM News 75 hosted a joint special session edited by Pedro Merino, coordinator of the WG, and Erwin Schoitsch, coordinator of the DES Dependable Embedded System WG, featuring almost 30 articles, many of which reporting advances on the application of formal methods in industry.

The FMICS workshop series has always been open to contributions from outside the ERCIM community, and strong links have been maintained with other organizations, such as Formal Methods Europe. In November 2009, the FMICS workshop will be held during the FM week, a special gathering of events organized this year by the FME association.

**Links:**

<http://www.inrialpes.fr/vasy/fmics>  
<http://www.win.tue.nl/fmweek>

**Please contact:**

Alessandro Fantechi – Chair FMICS WG  
 University of Florence and ISTI-CNR  
 E-mail: [fantechi@dsi.unifi.it](mailto:fantechi@dsi.unifi.it)

## Twenty Years of ERCIM News

*Twenty years of ERCIM stands also for twenty years of ERCIM News. ERCIM News was one of the first activities of ERCIM and with its 77 issues it has born witness to the advances in European R&D and Innovation in Information Technology and Applied Mathematics during the last twenty years.*

*“November, 1988: ‘Please come to CWI asap’. Agh, I was just sleeping off my jetlag. ‘Your colleagues await you in your office’, I was told upon arrival. There, Siggi Münch (GMD) and Laure Reinhart (INRIA) informed me that just now Baayen, Bensoussan, and Seegmüller were forging the ERCIM sword, and had ordered: ‘Present a plan for a newsletter at the plenary meeting, within three hours’. Our formula has - maybe surprisingly - not changed significantly over the years. The full support by our Directors was crucial in these first years, when ERCIM News was made with stone-age tools. It has remained great fun ever since, for me and – I am sure – for my colleagues as well.”*



Henk Nieland, ERCIM News editor 1989-2003



ERCIM News issues number 1 (1989) and 76 (2009).

“May this newsletter be one of our major channels through which we communicate with each other on our efforts, strategies, approaches and achievements, but also on our problems and questions”. That was the way the directors of CWI, INRIA and GMD, Cor Baayen, Alain Bensoussan and Gerhard Seegmüller started the first issue of the ‘ERCIM Newsletter’ in April 1989. Since then ERCIM News has evolved from an ‘in-house magazine’ to a publication covering reports and news about scientific projects from all over Europe and even beyond.

From the early issues on, ERCIM News has been published regularly four times a year. With each issue focusing on a special theme, the ERCIM News series has become a unique collection providing an overview on different topics of Information Technology.

ERCIM News was also a pioneer in online publishing. When the Web entered the research institutes in the early nineties, the question about an on-line version of ERCIM News was soon raised. From October 1994 on, the newsletter has been published both in printed and electronic format. The attractive full colour cover was first introduced in 1999, in concomitance with ERCIM’s tenth anniversary, and has been a crucial factor in the transition from an in-house newsletter to a popular scientific magazine.

Against all expectations, the electronic edition did not diminish the value of the printed edition but many readers on the Web subscribed to the printed edition, thus increasing its circulation. Beside the currently free subscription to the printed edition, the ERCIM News Web site gives access to all previous issues in electronic form.

ERCIM News today has a circulation of over 10,000 copies and is distributed in over 70 countries. More than 1500 articles have been published in twenty years, mainly written by scientists reporting on their research activities. The popularity ERCIM News has achieved is first due to the authors to whom the ERCIM Editorial Board wants to express their warmest thanks on this occasion.

**Link:** <http://ercim-news.ercim.org/>

## ERCIM Working Group 'User Interfaces for All'

*Constantine Stephanidis, Director of FORTH-ICS and Member of the ERCIM Executive Committee from 1998 - 2004 and member of ERCIM's Board of Directors since 2005; Chair of the ERCIM Working Group 'User Interfaces for All' (1995-2006) and Coordinator of the Coordination Action 'InterLink'.*



Active from 1995 to 2006, the ERCIM Working Group 'User Interfaces for All' had systematically promoted the proactive realisation of the Design for All principles in Human-Computer Interaction (HCI). The Working Group organised nine particularly successful truly international workshops, bringing together researchers from beyond different ERCIM institutions and other European academic and research organisations. The ERCIM UI4ALL Working Group was the recipient of the ERCIM WG Award in 2000. It constitutes a good practice example of how to successfully create and sustain a research community, and how to contribute to the establishment of a scientific field, in this case Universal Access and Design for All.

**Link:** <http://www.ui4all.gr/>

## European Coordination Actions: Beyond-The- Horizon & InterLink

Two Coordination Actions, administratively coordinated by ERCIM and scientifically coordinated by ICS-FORTH, have been funded by the FET Programme of the European Commission to set the bases for future research directions in critical ICT thematic areas. The first, named BEYOND-THE-HORIZON (from January 2005 to June 2006), addressed at a European level the thematic areas of Pervasive Computing and Communications, Nanoelectronics and Nanotechnology, Security, Dependability and Trust, Bio-ICT Synergies, Intelligent and Cognitive Systems, and Software Intensive Systems. The second, named InterLink (October 2006 - March 2009), addresses at an international level the thematic areas of Software Intensive Systems and new Computing Paradigms, Ambient Computing and Communication Environments, and Intelligent and Cognitive Systems.

**Links:**

<http://beyond-the-horizon.ics.forth.gr/>

<http://interlink.ics.forth.gr/>



*"I am very happy to celebrate together with the large ERCIM research community its 20th anniversary. Since its inception, ERCIM is pursuing the common vision of all European scientists for the creation of a truly European community of scholars, professionals and students in ICT and Mathematics; a forum where new young people can be trained, where new ideas can blossom and where collaboration and shared scientific goals can be achieved. It was an honor to participate in this great effort, as the second chairman of the Executive Committee and contribute to the common efforts during that early formative stage. ERCIM has my warmest wishes for even greater successes in the future and for remaining the vibrant organization that it is today."*

*Christos Nikolaou, ERCIM Executive Committee chair (1995 - 1998)*

*"Today, the importance of designing good interaction and involving users in our design processes is taken for granted. But when I did my PhD-studies, our work was considered a marginal slice of what computer science needed to care about. At SICS, we were a small group of four researchers, all PhD-students, who kept struggling with designing for end-users and worrying about design issues. In 1997 I was awarded the Cor Baayen award for my thesis. It made a huge change for me personally - it confirmed that what we had done was important and it opened many doors for me - but it was also important to the whole group at SICS. We moved from being a marginal group without any senior researchers, to a recognized and vital part of the institute. In 2000 we were already 25 people!"*

*Kristina Höök, 1997 Cor Baayen Award winner, leader of the Involve group at SICS and Professor at Stockholm University*



*"Spain joined ERCIM in 1993 through AEDIMA, the first 'consortium type' member of ERCIM but unfortunately this significant effort could not be continued in a stable way. Fast forward to 2003 when SpaRCIM, the Spanish Research Consortium for Informatics and Mathematics, was created as a more focused consortium representing Spanish computer science research in ERCIM. Over these years SpaRCIM has been sponsored by the Ministry of Science and Innovation (MICINN) and its decentralized organization is currently spread over five universities and one research institute from the Spanish Research Council (CSIC). SpaRCIM continues to promote high-quality research, participate in working groups, set-up joint projects, and disseminate research results."*

*Juan Jose Moreno-Navarro, Founding President of SpaRCIM*

## Eye to the Future – EchoGRID Roadmap: A Shared Vision for EU-China Cooperation on Next-Generation Grid Research

by Stephanie Parker

*EchoGRID (European & Chinese Cooperation on Grids), has focused on boosting interactions between researchers in Europe and China by creating gateways, defining common approaches, developing shared visions, and testing interoperability for standards development through a series of PlugTests.*

The final outcome is the EchoGRID Roadmap proposed for three, five, and ten year stages and addressing issues identified as being of mutual interest to researchers from academia and industry in China and Europe. New Programming Paradigms, Grid Architectures, Grid Management, Virtual Organizations, Component Models, Workflow – Business Processes are spotlighted as the top-level technological challenges for future collaborative developments in the drive towards the interoperability of open Grid services. Additionally, specific recommendations are targeted at EU and Chinese funding and research players, with the aim of facilitating the development of future research initiatives.

In order to make the EchoGRID vision a reality, there are a number of areas in which cooperation must be developed. Programming paradigm research activities must form the core of R&D projects; we must aim for dynamic service selection and business model support; the development of WS standards for Grids that also support backwards compatibility must be emphasized; and open-source competitors for tools that support distributed SOA infrastructures must be developed. Additionally, Grid management should place special emphasis on EU and Chinese research initiatives for SMEs and industry to experiment with innovative business models.

Mid-term goals for Virtual Organizations (VOs) include advancing their functionalities and subsequently aiming for semantic support in VOs. Further recommendations include the establishment of EU-China research groups to develop component models capable of supporting new properties of future Grid software; speeding up the development of programming frameworks for scientific and business applications; and ensuring the interoperability of different component models at different layers. Another important future goal is to define the composition of services as a focused research topic. Work in this area will bridge the gap between business and scientific workflows, while research projects bring the workflow/business process technology into sharp focus, with particular reference to far-reaching advances in the workflow tools.

A number of international conferences and focused workshops taking place in both regions over the last two years have served as important milestones for developing the roadmap. International events such as Beijing 2007 and 2008, Shenzhen 2008 and ICT2008 Networking Session have brought together enthusiasts and Grid experts from the Chinese Ministry of Science and Technology (MoST), the European Commission, EC-funded projects and Chinese and European enterprises. These events showcased successful international partnerships in Grids, software and services, and e-Infrastructures, and provided forums in which to deliberate top-level challenges and foster continued cooperation. Focused workshops have brought to the table high-profile experts from within and outside the project to define and make contributions to the main themes underpinning the roadmap. These events have also served to establish effective synergies with EU initiatives, such as BRIDGE, GRID-COMP, NESSI-GRID, RESERVOIR, 3S and Challengers, pinpointing areas of common interest and focus.

The EchoGRID project was funded by the European Commission from January 2007 to December 2008, and managed by ERCIM.

### Link:

The roadmap can be downloaded from:  
<http://echogrid.ercim.org/>

### Please contact:

Patricia Ho-Hune  
ERCIM office  
E-mail: [patricia.ho-hune@ercim.org](mailto:patricia.ho-hune@ercim.org)

## EuroIndia International Events

by Ashok Kar, Sudhanshu Rai and Florence Pesce

*The Euro-India ICT Cooperation Initiative (EuroIndia) has been holding a series of events to foster cooperation between Europe and India in various ICT domains targeted under the 7th Framework Programme. ERCIM is a partner in this project.*

EuroIndia organized a series of well-attended brainstorming and information days in India in 2008. The project also participated in the ICT 2008 event in Lyon, with a networking session on research and innovation cooperation between EU and Indian stakeholders across a series of ICT themes ranging from technology development to innovative applications. The information booth welcomed a wide range of delegates. Most sought information on ICT research and innovation activities and players in India, while some wished to explore concrete cooperation opportunities in well-defined ICT thematic areas.

January 2009 was an activity-packed month for the project, with three events organized in the last week of January in Mumbai and Delhi.

### Information Days

On 20 January, an information day was held in Mumbai at the prestigious Indian Institute of Technology with about fifty participants, the majority hailing from academia/research. The Delhi information day, organized on the sidelines of the first EuroIndia International Conference (see below) saw close to forty participants, equally divided between academia/research and other categories (industry, government, NGOs etc). The aim of these sessions was to give participants a comprehensive understanding of the 7th Framework Programme (2007-2013), the ICT Work programme (2009-2010) and the upcoming Call 4, as well as providing information on how Indians can engage in research projects alongside European partners.

### EuroIndia International Conference on ICT Innovations in India

The first EuroIndia International Conference took place on 22-23 January in New Delhi, in conjunction with the India R&D 2009: ICT Innovations conference organized by the Federation of Indian Chambers of Commerce and Industry

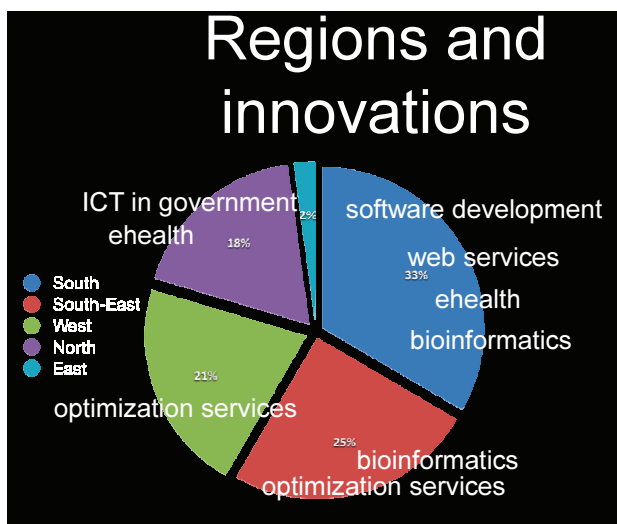


Figure 1: Share of publicly reported ICT innovation over 2003-2007 across regions of India. Leading domains of innovation in each region.

(FICCI). The EuroIndia ICT cooperation session was attended by about 150 delegates spread across academia (30%), industry (45%), policy makers and public research labs (20%), and media (5%).

EuroIndia made the most of a dedicated plenary session with Alvis Ancans from the International Relations Unit, DG Information Society & Media of the European Commission. In this session, a cross-section of players in Indian ICT research and innovation were given a complete picture of the ICT R&D organization, challenges, priorities and actions, instruments and funding.

The highlight of this session was a presentation of the findings of the first of four Indian ICT knowledge and innovation mapping modules by Sudhanshu Rai, team leader of this

mapping project and researcher in the Department of Informatics at Copenhagen Business School. The modules are:

- Historical
- Sample Survey
- In Depth Interview
- Delphi Trajectory Forecast.

Sudhanshu Rai shared with the audience some of the results of the first module, which has analysed publicly reported innovation in the ICT domains in India over the past five years across the five regions and within broad thematic categories. While some of the findings fit general expectations regarding the concentration of ICT R&D/innovative activity and domains, others were hitherto unreported (see Figure 1).

He then presented the other three modules of this mapping exercise, which will be rolled out during 2009. The findings of the mapping will be presented at the second EuroIndia International Conference on ICT Innovations in India, scheduled for November/December 2009.

### Upcoming Events

The next EuroIndia event will be a Brainstorming Day to be organized in Paris on 27 May 2009 in conjunction with ERCIM's 20th anniversary meetings. The two aims of this event will be to present the methodology and findings to date of the Indian ICT knowledge and innovation mapping that is currently underway, and to engage the EU ICT research constituency in order to synchronize EU and Indian ICT research and innovation priorities and policies. The agenda and registration will soon be available on the EuroIndia and ERCIM Web sites.

**Participation will be by invitation only. If you would like to contribute to this session, please contact: [info@euroindia-ict.org](mailto:info@euroindia-ict.org).**

The Euro-India ICT Cooperation Initiative ('EUROINDIA') is a Support Action under the Seventh Framework Programme coordinated by the Department of Informatics at the Copenhagen Business School. This initiative aims to reinforce collaborative research and innovation activities between Indians and Europeans in various ICT domains targeted under FP7. EuroIndia benefits from earlier groundbreaking work done under the BASIC (EuroIndia2004 Cooperation Forum on the Information Society) and MONSOON (<http://www.euroindia-it.org>) projects, which identified, profiled and engaged the widespread Indian ICT research community and helped increase Indian participation in FP6 and FP7 projects. ERCIM is a partner in the project.

**Link:**  
<http://www.euroindia-ict.org/>

**Please contact:**  
Florence Pesce  
ERCIM Office  
Tel: +33 4 92 38 50 73  
E-mail: [florence.pesce@ercim.org](mailto:florence.pesce@ercim.org)

Introduction to the Special Theme

# Future Internet Technology

coordinated by François Baccelli\* and Jon Crowcroft

The Internet has become critical to everyday life in domains as diverse as education, health, defence, commerce, travel and entertainment. The Internet was not designed for its current level of usage, and there is a need for simple constructs to allow the network to do better in terms of security, mobility and quality of service, among other things.

A number of global and local programmes of research (in the US, EU and further afield, eg Asia) are looking at future network architectures and building testbeds to evaluate new protocols and systems based on these new ideas. The most notable, in the US are the NSF Find programme and the GENI project to build infrastructure. In Europe we find a

number of projects, with some interesting high-level thinking (eg the Eiffel project) and a set of testbed initiatives under the FIRE Programme.

While some researchers think there is a need for a ‘clean slate’ design of the network of the future, much successful work is evolutionary rather than revolutionary, and this can be seen in the articles in this issue.

For this special theme on Future Internet Technology, we issued a call for papers and invited four well known researchers in addition, to cover the following topics as well as possible:

## Architectures and Infrastructures

Future core networks will leverage IP over simple, super-fast optical core networks. A major trend is that of virtualization, featuring the construction of optimized virtual networks that answer the needs of a collection of users or applications. A wide variety of wired and wireless access networks are available or being

*François Baccelli is a member of “GRIF” (Groupe de Réflexion sur l’Internet du Futur), together with Vivek Badrinath (Orange), François Bourdoncle (Exalead), Christophe Diot (Thomson), Serge Fdida (Université Paris 6) and Daniel Kofman (Télécom Paristech). The GRIF was launched by the French Minister of Industry.*

## Articles in this special theme

This special theme of ERCIM News contains 29 articles that are very much in line with the themes above:

### Setting The Context

- The invited article ‘**Starting the Debate: Agreeing on Disagreements**’ by Dirk Trossen, B.T., U.K., is a meta-discussion note, on how to trigger debate. As such, it is quite interesting as a methodology for encouraging architectural thinking in the long term.
- ‘**The Internet Engineering Task Force and The Future of Internet**’ by Emmanuel Baccelli, Thomas H. Clausen and Philippe Jacquet, INRIA Saclay, France, provides a useful reminder of the processes by which standards are created, and an analysis of how future research might be delivered in the short to medium term.

### Architecture

- ‘**Network Virtualization, a Perspective**’ by Anja Feldman, Mario Kind, Olaf Maennel, Georg Shaffrath and Christoph Werle: Virtualization, briefly defined above, can be seen as a way to overcome the ossification of the Internet, namely the general idea that one should not change a system or architecture that works well. The technical challenges and the business models associated with overlay networks and more generally with virtualization are thoroughly discussed in this invited paper.

- ‘**Y-COMM: A New Architecture for Mobile Heterogeneous Communications**’ by Glenford Mapp suggests that the symmetric architecture we are used to in the Internet is perhaps not well adapted to wireless, and proposes an alternative, useful for promoting discussion (cf the EU ANA Project).

### Modelling, Measurement and Management

- ‘**Epidemic Information Dissemination**’ by Laurent Massoulié (Invited paper): Epidemic algorithms became popular at the turn of the century with Peer-to-Peer applications. They are now being investigated by major companies designing efficient distributed live streaming mechanisms over the Internet. This beautiful line of research combines the design of innovative algorithms and the modelling of their execution over a large network, which allows one to prove their quasi-optimality.
- Understanding traffic entails delivering information efficiently to humans so that they can decide on actions to take. The article ‘**iMyNetScope- a Platform for Network Traffic Visualization and Analysis**’ by Pavel Minarik discusses one approach to this, and presents one particularly interesting topic for today - the element of Cyber Defence.
- ‘**Unified Access to Internet Measurement Data**’ by Felix Strohmeier, Martin Nilsson and

Demetres Antoniadis: like the previous article, this touches on a timely topic - making sure we can all (as researchers or operations) access data easily (see also [www.crowdad.org](http://www.crowdad.org) for an example of traffic archives in a standard form). Since the Internet is a federation of a vast number of networks (300 ISPs in the UK alone), we need clear and standard methods for exchanging data about operations. This applies equally to research networks such as those in FIRE.

- ‘**Privacy Aware Network Monitoring**’ by Peter Dorfinger, Carsten Schmoll and Felix Strohmeier describes a problem relevant to the research and operations community. Much work on the future Internet requires monitoring of research, but at the same time, real users require privacy. This article reminds us how important and difficult it is to reconcile these two aspects. Indeed, in gathering data about network use there are legal requirements relating to intercept laws and user privacy that must be met.
- ‘**Breaking the Weakest Link: Becoming a Trusted Authority on the Internet**’ by Marc Stevens is a reminder of the danger of security failures. The fact that even the basic technology of security is sometimes flawed is interesting. Systems are never perfect, and this is never more true than for security: we can never prove a system is secure, but only discover (eventually) when it no longer is!



developed: FTTH/FTTO in the wired domain, and 3G/CDMA, LTE, WiFi, WiMax and Satellite in the wireless setting.

### Internet Modelling, Simulation and Measurements

It is quite important to comprehend how the complex systems that we build actually operate. This requires basic research on network modelling and simulation, eg. in order to derive the fundamental laws on network dynamics and control or to evaluate the ultimate capacity of self-organized wireless networks. It also requires advances in network measurements: traffic statistics, Internet probing and measurement, network inference and detection of anomalies and attacks.

### Internet Algorithms and Software

The future of the Internet will require a wide range of computer science tools: verification; distributed algorithms (eg for consensus, and election and epidemic diffusion); resource management algorithms: resource allocation and scheduling; database algorithms: content storage, update and retrieval, content replication and consistency; search engines and the semantic Web.

### Self-Managing Networks

A key cost in networking is the operations and management overheads. The scalability of the Internet is well known, but as distributed applications proliferate, a more autonomic

approach is increasingly required. Spontaneous and self-organized networks emerge both in the wireless setting (eg with Wifi meshes or infrastructure-less wireless networks [MANETS]) and in the wired network setting (eg in peer to peer),

### Embedded Internet, Internet of Things

The Internet started by connecting computers and users to information, and then went on to connect users to each other with audio, video, games and social networking tools. Now the rapid evolution of pervasive, embedded networked devices means that we connect various types of devices to each other. There are far more computers in the world embedded in everyday objects (cars, domestic appliances etc) than there are on desktops. The interconnection of the Internet with the physical world through sensors and agents and the tagging of industrial production by RFIDs will lead to new traffic and architecture challenges, with possibly hundreds of billions of new devices that will collect information and will have to be upgraded and managed remotely and conveniently. This will require new paradigms for routing, search, naming, maintenance, data survival etc.

### Vehicular Internet

The Internet of things extends to cars and other vehicles. Here the key potential is in safety: ABS in cars could communicate road surface conditions to following vehicles, set-

### Correctness and Robustness (and self-managed to some extent)

- **‘ResiliNets: Resilient and Survivable Networks’** by David Hutchison and James P.G. Sterbenz: As discussed above, we need secure networks that are efficient and well managed. We also depend on these networks, and since individual components are never 100% immune from faults and failures, we need techniques to make the overall system more resilient than its components.
- **‘Standardized Testware for Internet-Based Telecommunication Services’** by Bostjan Pintar, Axel Rennoch, Peter Schmitting and Stephan Schulz: As Internet protocols grow in number and complexity each year, we need assurance that systems will behave correctly. This article outlines industrial solutions on the correct operation of the software systems (protocols) that make the Internet and its telecommunication infrastructure work.
- **‘Network Description Tools and Standards’** by Freek Dijkstra, Jeroen van der Ham and Ronald van der Pol: Setting up high-speed network connections - light paths - is still a manual effort taking two to three weeks. The network engineers need a clear picture of the network topology in order to plan and configure light paths. The Network Markup Language (NML) standardizes network topology and state information. The University of Amsterdam and

SARA Computing & Networking Services in the Netherlands are contributing to this effort, with the ultimate goal being to automatically set up and manage light paths.

### Novel Applications, especially multimedia

- **‘Delay-Tolerant Bulk Internet Transfers’** by Nikolaos Laoutaris and Pablo Rodriguez (invited paper): Using disruption-tolerant networking (DTN) techniques, but on the Internet rather than on challenged networks, produces very interesting results in terms of capacity. Throughput can be significantly increased for some applications, and there are some very well-motivated examples, such as backing up information between large data centres, using spare capacity in a smart and autonomic way.
- **‘Resource Management for IPTV Distribution’** by Henrik Abrahamsson and Per Kreuger: IPTV on AT&T and Telefonica’s networks now has in excess of 12 million users, each with hundreds of channels of broadcast quality content. This makes for non-trivial deployments, – and the importance of resource management for the various different technologies is unquestionable.
- **‘VISTO: Visual Storyboard for Web Video Browsing, Searching, and Indexing’** by Marco Pellegrini describes exciting work on future Internet killer applications.

- **‘The New Role of Humans in the Future Internet’** by Daniel Schall and Schahram Dustdar: Current service-oriented architectures typically orchestrate Web services. This article addresses architectures that combine Web and human services, where the latter are not limited to service consumption but may also undertake service production. The main emphasis is on the interfaces required in such architectures.

### In the Home (Internet of Things, and energy!)

- **‘Orchestrating your Surroundings’** by Pau Giner, Carlos Cetina, Joan Fons and Vicente Pelechano, UP Valencia, Spain, is an example of the ‘Internet of Things’ being at the centre of future Internet visions.
- **‘An Autonomic Home Networking Infrastructure’** by Thomas Luckenbach, Mario Schuster and Marc-Oliver Pahl: in line with the previous article but articulated with a vision of autonomic computing, this is on the important topic of home net automation. Home users are non-technical: they really do not want to become network managers, operators or system administrators. It is thus essential to make home networks ‘plug and play’.
- **‘Load-Balancing Energy-Usage of Household Appliances’** by Lennart E. Fahlen: Energy is a highly relevant topic, as described above, and using the Internet to control consumer

ting speed and braking reactions sooner rather than too late. Traffic management and pollution sensing on cars can also make use of networking between cars, and from cars to roadside infrastructure. Tracking goods in transit on the road would allow logistics companies to optimize their freight operations, saving time and energy and perishable goods.

#### Media Internet/Media-Driven Networks

There is an important diversification of the nature of the content transported across the Internet: initially it was files, then real-time games, video, telephony and whiteboards, and now TV, video on demand etc. There has also been diversification of the localization of contents: with each user potentially a content producer (peer-to-peer applications). New interactions with data are appearing, as in Web 2.0 or the semantic Web, and applications continue to evolve and require new systems, measurements and management tools.

#### Identity Management/Security

Since its beginnings in 1992, the World-Wide Web has offered remote transactions for goods and services. Recently we have seen a rapid growth in the number of attacks on identity, since acquiring such information allows miscreants to commit fraud that is hard to detect.

#### Energy saving and the Internet

Estimates vary, but the Internet and all its services consume something on the order of 4% of the energy in the developed

world. Only simple measures are required to improve this by a factor of two. Furthermore, the Internet, as we have discussed above, can be used to monitor and control external devices (things, vehicles, services) and significantly reduce their unnecessary power consumption. Figures as high as 30% have been quoted for possible national savings of energy, if unused devices in all homes could be remotely turned off. The investment necessary to achieve this is relatively low, with what seems like a very big potential return.

#### The Internet Services

We are moving from a Web of documents to a Web of services and Web of knowledge. This has triggered an explosion of new applications such as SecondLife, FaceBook and LinkedIn. And there is more to come, with augmented reality, virtual worlds, real-time games and telepresence. New concepts such as service orchestrations are also emerging, and with them, new business models.

#### Please contact:

François Baccelli

INRIA and École Normale Supérieure, France

E-mail: [Francois.Baccelli@ens.fr](mailto:Francois.Baccelli@ens.fr)

Jon Crowcroft

University of Cambridge, UK

E-mail: [Jon.Crowcroft@cl.cam.ac.uk](mailto:Jon.Crowcroft@cl.cam.ac.uk)

## Articles in this special theme

electronics in the household is a promising way to reduce unnecessary consumption.

#### On the Road, Vehicular

- **‘The Internet of Vehicles or The Second Generation of Telematic Services’** by Markus Miche and Thomas Michael Bohnert: Like the Internet of Things, vehicular networking is just taking off. There are car-to-car and car-to-infrastructure communications standards as well as integration of satellite navigation, traffic monitoring, tracking, and cellular map systems.
- **‘Web-Enabled Tracking Operations in Distributed Supply-Chains’** by Zsolt Kemény and Elisabeth Ilie-Zudor: Using the Web to track logistical information about goods in the real world is already a huge business. This is a very useful article on the topic, which also relates to the previous article.

#### Awareness, Identity and Society

- **‘Knowledge-Based Collaboration Patterns in Future Internet Enterprise Systems’** by Gregoris Mentzas and Keith Popplewell is supporting virtual community patterns, a bit like eScience community work.
- **‘Experiences from the Public Safety Arena for the Future Internet’** by Socrates Varakliotis,

Peter Kirstein and Steve Hailes: Since the Boxing Day tsunami and Hurricane Katrina, it has become clear that the Internet might be better used in support of communications during disasters. This is a useful note on work in the area.

- **‘With Joint Forces - Establishing Situation Awareness for the Future Internet’** by Tanja Zseby and Thomas Hirsch: Like Mentzas’ article, we can build systems to track patterns of use. This can be used to support collaboration or to detect misbehaviour. Increasingly, though, agencies are beginning to monitor social networks for this purpose.
- **‘Social Networking for a Pervasive Future Internet: the SOCIALNETS Project’** by Stuart M. Allen, Marco Conti, Andrea Passarella and Roger M. Whitaker: A timely article on an increasingly important topic. This outlines several new aspects of social networking that our readership should know about.
- **‘Owner-Centric Networking: A New Architecture for a Pollution-Free Internet’** by Claude Castelluccia and Mohamed Ali Kaafar proposes a new Owner-Centric Networking architecture that provides individuals control over their contents. This architecture would considerably improve privacy on the Internet by limiting data pollution.

- **‘Semantic Web for e-Commerce’** by Bernd Gruber: The Web is the basis for most new applications (Web mail, photo sharing, social nets, maps) and is the basis for commerce, which drives real economies.

#### Wireless, Embedded

- **‘Management of Future Mesh-Based Radio Access Networks’** by Vangelis Angelakis, Vasiliios Siris and Apostolos Traganitis: Community mesh wireless nets are becoming pretty widespread in some areas as an alternative to expensive cellular data services. Automating management is important; related work at Intel and Microsoft in the last couple of years has covered some aspects (AP selection) but the topic continues to be relevant.
- **‘Testing Mobile Data Applications on Smartphones with SymPA’** by Almudena Diaz and Pedro Merino discusses further the shortcomings of operating system platforms for data networking on cell phones and related devices. This is important, since the largest growth in Internet-capable devices is currently in this area.

# Starting the Debate: Agreeing on Disagreements

by Dirk Trossen

*The EIFFEL support action has been designed to mobilize European researchers to discuss and debate the future of the Internet, thereby helping to develop a functional networked society for the future.*

## EIFFEL: Created from the Need for Debate

Within the 7th Framework Programme, the EIFFEL specific support action (SSA) was created from the need to assemble recognized contributors to perform research on the future Internet in retreat-like meetings. As the project plan states: "...[the EIFFEL group] recommended that the best possible 'minds' (people) should be mobilized and motivated to establish European visions and to ensure proper interaction and debate to make the different vision trajectories meet in the future networked society."

EIFFEL will therefore gather international experts with the goal of developing visions for the future Internet. These meetings take place outside given funding structures: based on peer recognition the think tank members are invited to generously donate their time and knowledge. At the time of writing, the second of these think tank meetings had just finished, with some 25 external experts gathering in London in February 2009. To ensure lively debate and brainstorming, the number of participants at each meeting is kept relatively small. The invited members vary slightly, although it is intended that the EIFFEL community should gradually grow.

## ...And Not Alone in This Debate

The EIFFEL SSA is not the only initiative to undertake debate on the future Internet. Many European initiatives have been created over the past few years, most notably the Future Internet Assembly, as have initiatives in other regions of the world. The US initiatives on FIND (Future Internet Design) and GENI (Global Environment for Network Innovations) even precede European efforts like EIFFEL. For this reason, EIFFEL is attempting to attract a global audience, with international members being part of its think tank.

## Debating What Exactly?

Recognizing the need for debate is a relatively easy step compared to defining the subject of debate. It seems obvious

that an increased understanding of what we can commonly agree on should be the outcome; this was identified during the first two EIFFEL think tank meetings. But what is it that the community can agree on? Is it approaches for solutions towards the future Internet? Or common principles for a future architecture? All this is clearly desirable, but is the community at a stage where it can fulfill such demanding goals?

In order to agree on approaches that will be central to the future Internet, we

## Agreeing on Disagreements: Outlining an Approach

Pinpointing disagreements involves resolving contentions in some areas as much as formulating them in others. In other words, compiling a set of disagreements is more than merely listing every pet research topic and vision. We must determine why further dialogue is important, and this involves formulating potentially conflicting viewpoints on the importance of particular areas to the design of the future Internet. This does not oppose Alan Kay's view that "the

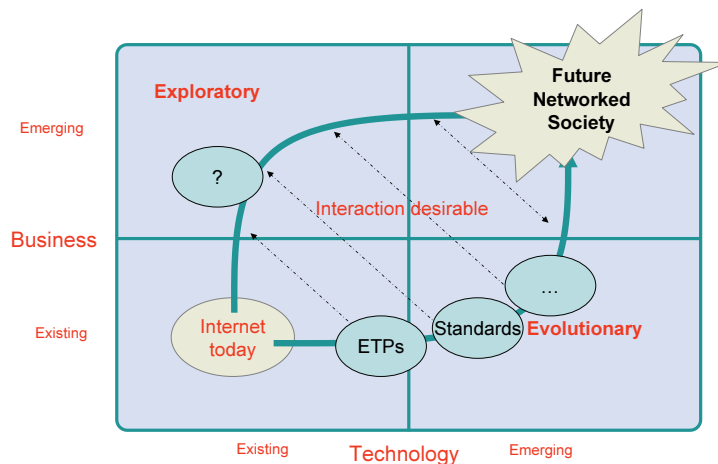


Figure 1: Need for Debate: A balanced approach.

as a community need to identify the contentious areas of investigation. Hence, we need to 'agree on the disagreements', which will lead us to the (potentially) necessary further investigation and debate around these disagreements. This debate will need to evaluate relevant solutions, hopefully paving the way for a growing set of agreements that could form the foundation for the future Internet.

The London think tank meeting saw the beginning of this debate, with presentations as well as break-out groups in areas like trust and security (David Clark), measurements required for future Internet research (Kimberly Claffy), the issue of inherent (social and moral) values for the future Internet (Ian Brown), and the operational viewpoints of a large data centre provider (Vijay Gillis).

best way to predict the future is to invent it", but complements it with the necessary debate on why that future is worthwhile pursuing.

Individuals must therefore approach the meetings ready to discuss and even defend their viewpoints with respect to the importance of the topics at hand. There must also be a component of community contribution, building on and recognizing the collective intelligence of the wider community. Overall, contributions to this debate, its disagreements and its potentially controversial viewpoints need to come from a community of researchers that feels strongly about the importance and relevance of its work to the future Internet.

However, some form of editorial steering of the debate is also necessary. This

will help consolidate debate in areas where consensus is beginning to emerge, while fostering and moderating it in areas where disagreement prevails.

Careful consideration of these issues has led us to consider a Wikipedia-like online platform as the first step for this debate. This platform is charged with gathering opinions and consolidating the debate of a growing community of researchers, and will be a repository for relevant work in areas that are critical to the making of the future Internet. The platform will most importantly serve as a tool with which to formulate, debate and finally resolve disagreements. The use of collaborative editing tools like wikis will allow this debate to be

recorded and structured. It is important to note that this online platform will differ from a conventional Wikipedia-like encyclopedia in having a clearly defined editorial team, soliciting contributions from identified (registered) users (while being open to be read by anybody), and allowing individual contributors to retain the copyright of their contributions.

At the time of writing, an initial team is in the process of setting up such a platform. It is expected to launch during the spring of 2009, seeded with results from the initial EIFFEL think tank meetings but also incorporating material from experts worldwide. While a particular timeframe will not be

enforced, we do expect the debate to pick up quickly, leading to early results.

The positions presented in this paper are the outcome of discussions among EIFFEL think tank members and therefore includes the viewpoints of many individuals involved in these discussions.

#### Links:

<http://www.fp7-eiffel.eu>  
 US initiatives FIND and GENI:  
<http://www.nets-find.net/>  
<http://www.geni.net>

#### Please contact:

Dirk Trossen  
 BT Research, UK  
 E-mail: [dirk.trossen@bt.com](mailto:dirk.trossen@bt.com)

## The Internet Engineering Task Force and the Future of the Internet

by Emmanuel Baccelli, Thomas H. Clausen and Philippe Jacquet

*The Internet Engineering Task Force was the birthplace of today's Internet. Understanding its activities is necessary for individuals and institutions who wish to anticipate the future of the Internet. As things stand, this necessity is not likely to fade any time soon.*

For those with a stake in the ideas and initiatives that will drive the Internet in the future, the Internet Engineering Task Force (IETF) is unavoidable. Created in 1986 by US government agencies (DoD, Department of Energy, NASA, NSF) to supervise the design and deployment of Internet protocols, it was initially open only to US government-funded researchers. Early 1987 saw a dozen industry representatives invited, and in a matter of months, the IETF was opened to all interested parties. In 2008, IETF meetings were attended by roughly 1300 engineers and researchers from all over the world.

The IETF is an R&D forum in which network engineers define, describe, review and discuss network protocols, which are published as Requests For Comments (RFC). These then may or may not be implemented and used by industry. IETF meetings are triannual, with business in the interim being conducted on open mailing lists.

#### Organizational Structure of the IETF

Work within the IETF is organized into working groups (WGs), each of which is in charge of a specific problem (eg

mobile ad hoc routing). Typically, a WG is supervised by two chairs.

WGs within the same general field are assembled in a so-called 'area' (eg the routing area). Each area is supervised by two area directors (AD), whose task is to shepherd the creation, activity and eventual demise of WGs in the area. In early 2009, the IETF had eight areas and 120 working groups.

The assembly of area directors forms the Internet Engineering Steering Group (IESG). The IESG, together with the Internet Architecture Board (IAB), ensures the overall coherence of the Internet protocols 'corpus'. IESG and IAB members are periodically replaced, potentially by any other competent IETF participant.

#### The IETF and Decision Making

In contrast to standardization bodies such as IEEE or ETSI, individuals represent themselves to the IETF: there is no de facto company representation. People from the same company/institution may make conflicting contributions, while people from different com-

panies/institutions may contribute together to a standard without the necessity for a formal agreement. Proposals must be open for other potential contributors without any copyright restrictions. Moreover, the IETF's fundamental motto is: "We reject kings, presidents and voting. We believe in rough consensus and running code."

Rough consensus: rather than voting (as in the IEEE or ETSI), decisions in the IETF are made based on 'rough consensus'. In a WG this is gauged by the WG chairs, and in the IETF as a whole it is gauged by the IESG. Well understood by IETF participants, this procedure allows any good idea from any origin to be discussed, bringing contributions from individuals and small institutions on an equal footing with those from big companies.

Working code: generally a proposed protocol cannot be promoted as a potential standard without thorough experimentation. Experiments on protocols can be performed using working code and minimal hardware investment, often none. Furthermore, to avoid artefacts due to internal bugs, several working-

code bases developed independently following the proposed specification must demonstrate their full compatibility before the standard can be validated.

#### The Pertinence of the IETF

The ability of an R&D forum to meet the positive evolution of a technology depends on how it manages the four following parameters: vision, legacy, luck and necessity.

*Vision:* the IETF clearly has the right focus. While its vision is fuzzy, since initiatives generally come from the bottom, its top-level directions are very clear. Currently, for instance: mobility, scalability to encompass the Internet of objects, or IPv6. Introduced in the 1990s to address the scarcity of available addresses with IPv4 (four bytes format), IPv6 upgrades IP to a flexible address management scheme over 16 bytes, potentially identifying 1038 elements. While transition from IPv4 to IPv6 is slower than expected due to the generalization of CIDR, NAT, and DHCP, experts predict the allocation of the last IPv4 address to take place in 2010.

*Legacy:* an R&D forum is the meeting place for dreams and possibilities. However, the most brilliant idea in the world may be presented in vain if it is incompatible with existing technology:

“A good idea is not always a good idea”. Nevertheless, the IETF is very careful not to bypass any innovative idea, and manages to this end a parallel forum called the Internet Research Task Force (IRTF), where new paradigms (eg delay-tolerant networking) are trained to fit legacy.

*Luck:* the most important issue in an R&D forum is the ability to manage an unexpected breakthrough. With a culture of ideas beginning at the bottom, even the most crazy idea is welcomed if it ?ts legacy and addresses a concrete issue: “A good idea can become an extremely good idea”. A striking example is TCP. In the late 1980s, the challenge was to cope with brutal capacity reduction when data traffic had to cross long-haul networks. Failing to address this issue caused the demise of a concurrent system, ATM. The IETF produced a surprisingly simple, but innovative, solution: with TCP, a source terminal tunes the ?le transmission pace according to feedback from the destination terminal. Experts consider the strength of TCP (supporting variations of network capacity ranging over more than twelve orders of magnitude) to be the main reason for the success of the Internet.

*Necessity:* the IETF mandates itself to solve certain problems. For example in

the late 1980s, the current routing protocol RIP failed when a set of routers was brutally removed from the network. This bug, called ‘count to infinity’, created a sustained loop that caused an avalanche of disruptions: the Internet was down for two full days. A failure indeed for a system designed with resilience as its core tenet! RIP had to be replaced by a new protocol, specified in emergency: Open Shortest Path First (OSPF), widely used nowadays. Less elegant than RIP, OSPF is far more robust, based on an exhaustive mapping of network links that allows routers to compute new routes and react in real time to disruptive topology changes.

#### The IETF and the Future

The IETF was the birthplace of the Internet of today. Understanding its activities is necessary for individuals and institutions who wish to anticipate the future of the Internet. As it appears, this necessity is not likely to fade any time soon.

#### Link:

<http://www.ietf.org>

#### Please contact:

Emmanuel Baccelli

INRIA, France

Tel: +33 169334101

E-mail: [Emmanuel.Baccelli@inria.fr](mailto:Emmanuel.Baccelli@inria.fr)

#### Invited Article

## Network Virtualization - An Enabler for Overcoming Ossification

by Anja Feldmann, Mario Kind, Olaf Maennel, Gregor Schaffrath and Christoph Werle

*While the Internet is currently viewed as widely successful for some of its participants, namely the users and service providers such as Google, it also suffers from ossification in the underlying infrastructure. The ossification has multiple causes, among them the fact that since the Internet works quite well as it is, Internet Service Providers (ISPs) have no incentive to change their ways. Moreover, ISPs suffer from a lack of business perspective due to the predominant charging modi for Internet access: flat rates for users and a combined price model consisting of a base rate and usage-based component for content providers. An additional complication is that traffic grows at a higher rate than that at which the network equipment costs decrease.*

As a consequence, there is debate within the ISP community over whether ISPs should become pure bit-pipe providers or should offer value-added services. In addition, some service providers have found that some application support inside the network can help their applications and are considering investing in network infrastructure, eg Google.

Indeed, over the last twenty years almost all innovation, eg novel applications, has taken place at the edge of the network, while the core remains almost untouched. However, the time has come to support novel applications with services inside the network (eg via network-based enablers), and to revisit the Internet architecture to add native sup-

port for security, mobility and manageability.

For circumventing the difficulty of changing successful networks, the concept of overlays has proven to be very useful. For example, the Internet got started as an overlay on top of the phone network. One of the key insights is that

each overlay can be considered a virtual network.

Virtualization is an old but very successful technique for CPU, memory, storage and almost all other system resources. Fundamentally, virtualization is an abstraction concept that hides hardware details, eg to cope with heterogeneity. It effectively offers a level of indirection as well as resource sharing. The former improves flexibility while the latter enables partitioning as well as reuse of resources, resulting in higher efficiency. However, to achieve this it also requires a resource separation with a sufficient level of isolation.

Within the last five years, end-system virtualization, eg via Xen or VMware, has revamped server business. Router vendors such as Cisco and Juniper offer router virtualization, and existing techniques such as MPLS (Multiprotocol Label Switching), GMPLS (Generalized MPLS) and VPNs (Virtual Private Networks) offer some coarse-grained link virtualization. Overlays such as peer-to-peer (P2P) networks over the Internet (eg BitTorrent) can also be seen as a virtual network, but they suffer from a lack of sufficient isolation. VPNs (eg realized via MPLS), can also be seen as virtual networks. However, they suffer from a lack of node programmability.

Indeed, a significant part of the current Internet infrastructure either already supports or has the potential to support a basic form of network virtualization. Moreover, due to high operational costs, some sharing of network resources among network operators already exists. For example, T-Mobile UK and 3UK share network sites. As such we need to explore the technical feasibility and potential business opportunities that virtualization can offer while overcoming Internet ossification.

The first observation is that it is possible to treat the current Internet as one future virtual network, which implies that one does not have to 'change the running system'. The next observation is that service providers can potentially operate their own virtual network according to their needs, eg, to offer a value-added service. A virtual network here may imply operating a non-IP network that may require low-level access to each 'slice' of each network device.

We also point out that each of these virtual networks can be built and operated according to different design criteria; for instance, they could optimize a specific network metric like throughput, latency, or security. This is possible as long as the virtual networks are properly isolated, which requires (among other things) corresponding Quality-of-Service (QoS) support in the underlying network. For example, one network might be optimized for anonymity while another is optimized for accountability. Virtual networks offer the added benefit to the service provider that their resources (eg node or link resources as well as the topology) can be increased or decreased gradually in line with the popularity of the service. They offer the benefit of resource migration and resource aggregation to the network infrastructure provider.

Nevertheless, network virtualization creates a tussle among service and infrastructure providers over who should operate and who should manage such virtual networks. Hence, there is a need for additional players besides providers (PIPs) and service providers (SPs): virtual network providers (VNPs) for assembling virtual resources from one or multiple PIPs into a virtual network, and virtual network operators (VNOs) for the installation and operation of the VNet provided by the VNP according to the needs of the SP.

In terms of business relationships, the VNP buys its bit pipes from one or several PIPs. The VNO uses the resources assembled by the VNP to operate a virtual network according to the needs of an SP. Note that virtualization enables a VNP to act as a PIP to another VNP.

#### Requirements and Conclusion

To realize the benefits of virtualization, we need an architecture for network virtualization that encompasses the players, PIP, VNP, VNO and SP. On the technical side, we need standardized interfaces between the players to automate the setup of virtual networks, ie, a common control plane. Moreover, we need ways in which each player can check if it is being provided with the service it is paying for (eg in terms of QoS). Furthermore, it must be possible for the PIP to render/delegate low-level management of the virtualized pieces of the network components via the VNP to

the VNO. Of course for virtualization to succeed it must be accepted by each player that some information is hidden, for example the SP should not be able to know exactly which link within the PIP is being used by a certain connection. However, not all information can be hidden; for example, an SP might want to specify the reach of its virtual network. Therefore, we need to explore the trade-offs between the level of specification of the virtual network and the flexibility of optimizing the resource usage. Another challenge associated with information hiding is 'debuggability', eg, figuring out who is responsible if something does not work as it is supposed to. Still, the unique opportunity is that when business processes are being redesigned, it will be possible to completely trace a link in a virtual network as used by the SP to a link in a PIP. On the business side, we need agreements on the kinds of contract and their prices. Moreover, we need support for accountability, audits and security.

Virtualization shows a lot of promise. Nevertheless, each player must determine for itself the role it wants to take. For this, it must analyse its business position and determine the potential benefits and dangers of each role. As such, further work is needed on both the technical and the business aspects of network virtualization.

#### Link:

4WARD Project:  
<http://www.4ward-project.eu>

#### Please contact:

Anja Feldmann, Olaf Maennel,  
Gregor Schaffrath  
TU Berlin/Deutsche Telekom  
Laboratories, Germany  
E-mail: {anja, olaf, grsch}@net.t-  
labs.tu-berlin.de

Mario Kind  
Deutsche Telekom Laboratories,  
Germany  
E-mail: mario.kind@telekom.de

Christoph Werle  
Karlsruhe University, Germany  
E-mail: werle@tm.uka.de

# Y-COMM: A New Architecture for Mobile Heterogeneous Communications

by Glenford Mapp

*Y-Comm is a new communications architecture that will meet the challenge of providing ubiquitous heterogeneous communication on a global scale. This new Internet will provide continuous connectivity by the seamless operation of multiple mobile networks that will simultaneously be accessible by mobile nodes. It will also provide transparent support for quality of service (QoS), fostering the development of new kinds of applications. Finally, it will provide built-in multi-layer security.*

Y-Comm is based on the observation that the widespread deployment of mobile wireless networks represents a key evolutionary step in the development of the Internet. It splits the Internet into two distinct components: a core network comprising an optically switched backbone surrounded by fast access networks, and peripheral networks dominated by different wireless technologies. Y-Comm seeks to represent this new reality as an architecture comprising two frameworks. The first, the Peripheral Framework, runs on the mobile node and interacts with peripheral wireless networks. The second, the Core Framework, runs in a distributed fashion in the core infrastructure. The Y-Comm architecture is shown in Figure 1.

In addition to the Core and Peripheral Frameworks, Y-Comm has a multi-layer security system that works with the two frameworks to provide a completely secure environment.

## The History of Y-Comm

Y-Comm came about because of work done on the Cambridge Wireless Testbed at the University of Cambridge. Designed for studying vertical handover, this testbed was built using several networks including WLANs, GPRS and LANs. End devices consisted of desktop machines doing handovers between GPRS and local LAN as well as mobile devices doing handover between WLAN and LAN systems.

The Cambridge Wireless Testbed built mechanisms to implement client-based vertical handover in which the mobile node, not the network, controlled the handover. The mobile node monitors the state of its various network interfaces and decides when and where handover should occur. In addition to client-based handover, a reactive policy mechanism called PROTON was invented to look at when handover should occur.

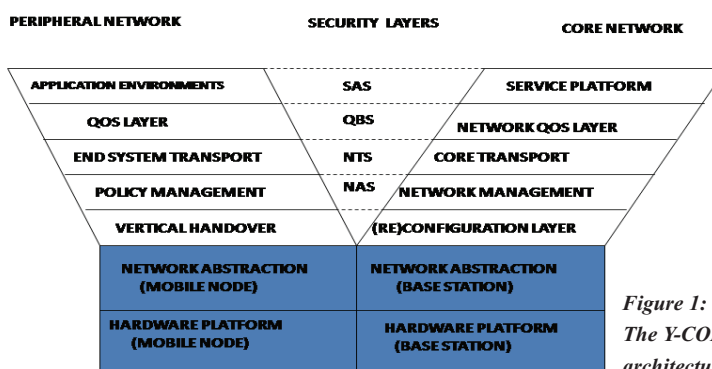


Figure 1:  
The Y-COMM  
architecture.

Following this effort, the key issue was this: how would you build a commercial network to provide seamless ubiquitous connectivity? It soon became clear that a new framework was needed to build future mobile systems. Traditional frameworks such as the OSI model are inadequate to meet these challenges, and so Y-Comm was born.

## The Y-Comm Research Group

The Y-Comm research group comprises a number of academic institutions exploring a variety of aspects of Y-Comm design. It includes the School of Engineering and Information Sciences (EIS) at Middlesex University in London, the Computer Laboratory at the University of Cambridge, and The Institute of Science, Mathematics and Computing at the University of Sao Paulo in Brazil. Other institutions include Loughborough University, the University of Calgary, Dartmouth University USA and the University of Trinidad and Tobago.

## Developing New Technologies for the Internet of the Future

One of the great outcomes of studying the Y-Comm architecture is that it has helped us to clearly define the new technologies that are needed to support future heterogeneous mobile networks. For the Peripheral Framework, this includes cognitive radio, vertical handover and policy management mechanisms as well as QoS-aware middle-

ware systems and the ability to build new application environments using these underlying technologies. For the Core Framework we need programmable techniques to manage key parts of the mobile infrastructure, network virtualization techniques to manage networks in different geographical areas, core-endpoint management, and service platforms for the widespread deployment of network services and applications. A new research network is being formed to develop and bring together these diverse technologies.

## The Y-Comm Testbed

Over the next few months, we hope to begin work on a Y-Comm testbed, which will be used to test algorithms and mechanisms needed to implement the Y-Comm architecture. The testbed will initially support vertical handover between WLAN, 3G and LAN systems and will first concentrate on demonstrating a complete implementation of the Peripheral Framework running on a mobile node using the Android Operating System.

## Link:

<http://www.mdx.ac.uk/eis/research/groups/YComm.asp>

## Please contact:

Glenford Mapp  
Middlesex University, UK  
Tel: +44 208 411 6871  
E-mail: [G.Mapp@mdx.ac.uk](mailto:G.Mapp@mdx.ac.uk)

# Epidemic Information Dissemination

by Laurent Massoulié

*At the Thomson Research Lab in Paris, epidemic algorithms are being investigated. These algorithms are expected to have a strong impact on the dissemination of media content over the Internet, notably in peer-to-peer systems for live streaming applications.*

While epidemics are part of everyday life, they nonetheless constitute an intriguing phenomenon. In particular, one may wonder at how quickly a global outbreak (say of influenza) occurs, given the local and limited nature (individual contacts) of the propagation mechanisms of the infectious agent.

The earliest attempt at a mathematical understanding of these dynamics can be traced back to the Swiss scientist Daniel Bernoulli (1700-1782) in his study of smallpox in 1766. A subsequent mile-

stone was reached in 1838 by the Belgian mathematician Pierre Francois Verhulst, who described the temporal evolution of a population size by the so-called logistic equation (see box). The most salient feature of this equation is its exponential growth in the early stages.

for this by inventing a way of putting epidemic propagation to good use. Specifically, they have solved the delicate problem of synchronizing distributed databases by spreading data updates among individual databases by random, viral-like propagation attempts.

This seminal work introduced the idea of mimicking biological epidemics to propagate useful information: the area of epidemic algorithms was born. In a nutshell, this approach lets each con-

peers. This data stream is fragmented into data chunks by the data source. Peers subsequently exchange data chunks among themselves to eventually recover the original data stream.

An epidemic scheme for this live streaming application is then specified by detailing how peers choose which of their neighbours to contact, and which data chunk to forward to the selected peer. As a result, after appearing at the data source, each chunk should spread within the system to reach all peers within a limited delay. However, the epidemics corresponding to individual chunks compete for communication resources. The challenge is then to define mechanisms for which this competition does not slow down individual epidemic disseminations. We now consider several scenarios for which we have identified efficient epidemic schemes.

Symmetric networks: Assume all peers have the same communication capacity of one chunk per second. Then the time for a chunk to reach all peers in a system of  $N$  such nodes is at least  $\log_2(N)$ , since in one second the number of recipients of a chunk can at most double. Moreover, the rate at which peers can all receive data cannot exceed one chunk per second. In this context, consider the Random Peer Latest Useful Chunk scheme, whereby each peer selects a target at random, and sends to the elected target the latest chunk it received, that the target has not yet received.

Recently we (T. Bonald, L. Massoulié, F. Mathieu, D. Perino and A. Twigg) established that for any data rate strictly below 1, this scheme ensures that each peer receives chunks within the optimal time  $\log_2(N)$ , plus some constant (not depending on  $N$ ). Thus this scheme achieves delivery in optimal time for any injection rate strictly below 1.

Our proof builds upon previous work (by S. Sanghavi, B. Hajek and L. Mas-

## Logistic equation



*The Belgian mathematician Pierre Francois Verhulst described the temporal evolution of a population size by the so-called logistic equation:*

$$f(t) = \frac{f(0) \times e^{\beta t}}{1 - f(0) + f(0) \times e^{\beta t}}$$

stone was reached in 1838 by the Belgian mathematician Pierre Francois Verhulst, who described the temporal evolution of a population size by the so-called logistic equation (see box). The most salient feature of this equation is its exponential growth in the early stages.

With the advent of personal computers, the eighties also witnessed the invention of computer viruses. Their propagation is as fast as that of biological epidemics, and is accurately modelled by the logistic equation of Verhulst. (This has been strikingly illustrated by the latest incarnations of computer viruses, namely Internet worms, and in particular the so-called Code-RED worm.) Clearly, this is not an invention that computer scientists should be proud of! In the same period however, a group of computer scientists from Xerox PARC have compensated

stituent of the system choose at random which of its neighbours to 'infect' with information. This is certainly simple, but is it efficient?

Surprisingly, in a number of situations of interest, epidemic approaches are as efficient as any other technique, no matter how elaborate. Let us illustrate this in the context of peer-to-peer systems. These consist in collections of machines (the peers) which act both as servers and as clients of services, and interact over the Internet. Such systems became popular for file-sharing applications with Napster at the turn of the century.

More recent versions of these systems provide so-called live streaming applications: a stream of audio-video data (think of a TV channel) is to be disseminated in real time to all participating



soulié) where the performance of the Random Peer Latest Blind Chunk policy is analysed. The argument in this analysis relies on the identification of secondary epidemic disseminations, which are shown to follow a variant of the logistic equation.

Heterogeneous networks: Assume now that peer uplink capacities are no longer identical. In this context, we considered the Most Deprived Peer Random Useful Chunk policy, whereby each peer picks the neighbour to which it can provide the largest number of chunks, after which it selects at random which of these chunks to forward. We (L. Massoulié and A. Twigg) showed that this scheme ensures delivery of all chunks in bounded time, provided the data rate

is strictly below the optimal injection rate. In this sense, this is a rate-optimal scheme.

The above results show that epidemic schemes can deliver live streaming at optimal performance in several scenarios of interest. Thus randomized local decisions can lead to globally optimal behaviour. Peer-to-Peer architectures are expected to become the principal channel for delivering media content (legal or otherwise) in the near future. Epidemic algorithms will then have a strong impact in driving the dissemination of all such media content.

Many exciting problems remain to be solved in this area: suitable schemes for delivery of video-on-demand are being

studied; delay optimality in heterogeneous environments is still poorly understood; and the search for schemes to ensure that data flows reduce the cost on the underlying networks (the so-called ISP-friendliness issue of peer-to-peer design) is also wide open.

**Link:**

The above-mentioned articles can be found at <http://www.thlab.net/~lmassoul/>

**Please contact:**

Laurent Massoulié  
Paris Research Lab, Thomson,  
Corporate Research, France  
E-mail: [laurent.massoulie@thomson.net](mailto:laurent.massoulie@thomson.net)

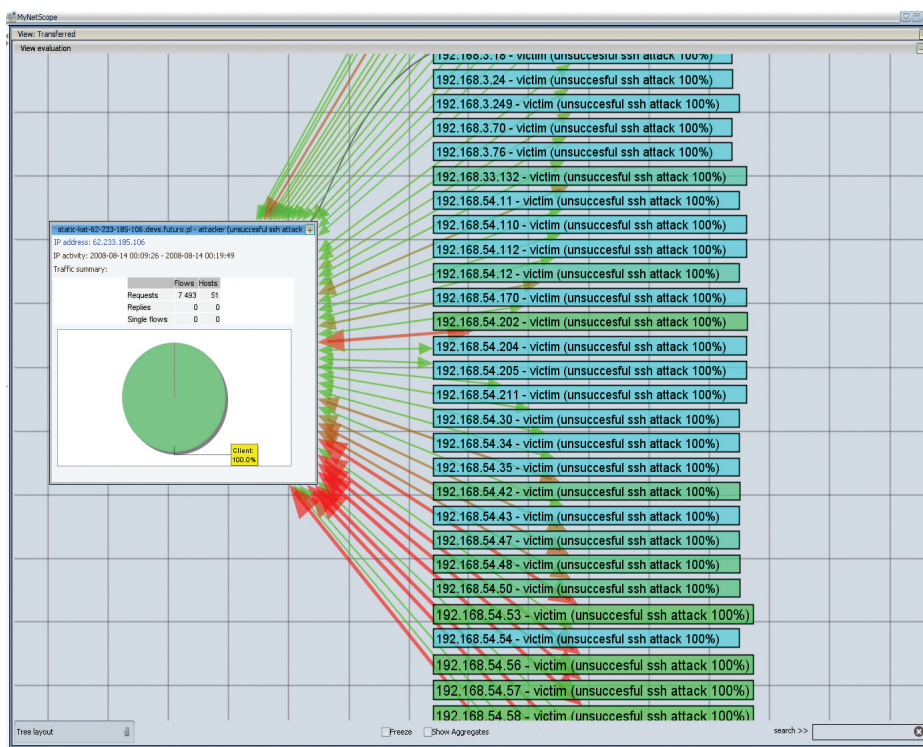
## MyNetScope: A Platform for Network Traffic Visualization and Analysis

by Pavel Minařík

*The MyNetScope development project is an example of academic and business cooperation resulting in a platform for advanced network traffic processing, analysis and visualization. MyNetScope overcomes the barrier of traffic content by focusing on traffic characteristics and behaviour patterns and targets the intrusion detection and prevention systems (IDS/IPS) segment of this century.*

The MyNetScope platform is a reaction to the current tendency to analyse or process network traffic using statistical methods. While traffic content processing is inapplicable in encrypted traffic, statistical methods are unable to detect precisely targeted or sophisticated attacks. The MyNetScope platform overcomes these limits by performing behaviour-based analysis. It moves down from the application layer to the network and transport layer, reducing the amount of data to be processed but still focusing on individual data flows. Typical tasks for behaviour analysis include dictionary attacks against network services or the misuse of secured hypertext transport protocol (HTTPS), where signatures can't be specified and a statistical approach may detect only massive attacks or HTTPS protocol misuse.

MyNetScope focuses on data flow characteristics and statistics – NetFlow data – and doesn't work with traffic content. NetFlow is an open industry standard defined by Cisco and designed to monitor large-scale and high-speed networks.



*Figure 1: The results of SSH dictionary attack detection on the university network (university IP addresses were obfuscated). This attack took place repeatedly during August 2008. The expanded graph node on the left represents the attacker and some attack details, while the nodes on the right represent victims. The probability of attack success is also available.*

MyNetScope provides an interactive insight into network traffic. It combines various visualization methods (dynamic mind maps, tables, forms and statistical graphs) in a single workspace and guides the user through visualization showing a greater or lesser level of detail according to the user's preference.

The main purpose of MyNetScope is the detection of anomalies based on behaviour patterns. Behaviour patterns are built upon individual NetFlow data processing and IP address classifications. MyNetScope is able to distinguish between requests and replies (servers and clients) in the network since it implements the Bidirectional Flows Standard as one of the first systems. A successful example of behaviour pattern is dictionary attack against SSH (secured shell) service which was developed at the Institute of Computer Science of Masaryk University. MyNetScope is open to extensions, meaning it is possible to add new detection methods to the application. MyNetScope is also able to integrate data from various sources, eg, it uses DShield portal to support detection

methods and identify world-famous attackers.

The MyNetScope development project was started in September 2007 by a Masaryk University spin-off company called Mycroft Mind. Since January 2008 it has run in cooperation with the Institute of Computer Science at Masaryk University, combining the technology provided by the company (visualization engine, memory engine, pattern-matching engine and data integration engine) with expert knowledge from university. Its alpha version was finished in April 2008 and deployed on the university network of Masaryk University in Brno, which comprises about 10,000 computers, with more than 200 million data flows per day. In January 2009 a beta version was deployed and is being used to detect dictionary attacks against the SSH service and in the investigation of security incidents.

MyNetScope is also being used in other research activities, eg, grant Cyber defence of the Czech Ministry of Defence, the goal of which is to provide the knowledge to face organization-spe-

cific and other new types of threats. The MyNetScope development team currently focuses on the detection of network infiltration or misuse of HTTP and HTTPS services (tunnelling using these services) and its detection is based on NetFlow data characteristics.

While the application is heading to the domain of IDS/IPS, it focuses on behaviour detection and specific threats which are out of reach of common IDS/IPS. Another advantage is the ability to visualize network traffic and explore suspicious behaviour on the network. The project goal from the point of view of Masaryk University is to transfer the research and development results into industry.

**Link:**

MyNetScope:  
<http://www.advaict.com/mynetscope>

**Please contact:**

Pavel Minařík  
Institute of Computer Science, Masaryk University/CRCIM  
Tel: +420 605 852 335  
E-mail: [minarik@ics.muni.cz](mailto:minarik@ics.muni.cz)

## Unified Access to Internet Measurement Data

by Felix Strohmeier, Martin Nilsson and Demetres Antoniadis

*Optimizing control, management and flexibility of future network infrastructures requires a good understanding of network and application behaviour. The FP7 project MOMENT (Monitoring and Measurement in the Next-Generation Technologies) combines a diversity of data sources, enriches them with semantic information, and presents them via a unified interface in a user-friendly way.*

Since the Internet does not by design integrate network measurement or quality-of-service assurance mechanisms, it has always been necessary to perform measurements independently from the actual communication in order to obtain status information from the network. Deployed monitoring and measurement infrastructure and tools represent essential technology for migrating to a future Internet. Besides the per-domain operational monitoring performed by Internet Service Providers, research institutes and projects are establishing monitoring infrastructures for a large-scale view of the Internet. Owners of such infrastructures are willing to share their collected data for research and development, but usually only provide proprietary interfaces

such as Web-based graphical user interfaces, which do not allow programmatic subsequent processing of measurement results.

### Heterogeneous Data Sources

As depicted in Figure 1, the MOMENT project targets the creation of a 'mediator' that allows unified access to several current and future sources of network monitoring and measurement data. Existing measurement infrastructures and tools like Dimes and Etomic make their data available via different access methods. Three examples of data sources created by ERCIM members cooperating in the MOMENT project are presented in the following. First, MINER is the "Measurement Infrastructure for NET-

work Research" that orchestrates multiple measurement tools. It allows measurement scenarios to be specified and scheduled, and results to be retrieved from them. These services are provided by a tool-agnostic unified programming interface. Secondly, BART (Bandwidth Available in Real-Time) is a tool for measuring end-to-end available bandwidth in real time, generating a continuous stream of measurement data. Finally, Appmon is a traffic classification application able to attribute network traffic to the application of origin, by inspecting the payload of the packets. Appmon provides both real-time and historical usage patterns for a number of popular applications, such as BitTorrent, eDonkey and HTTP.

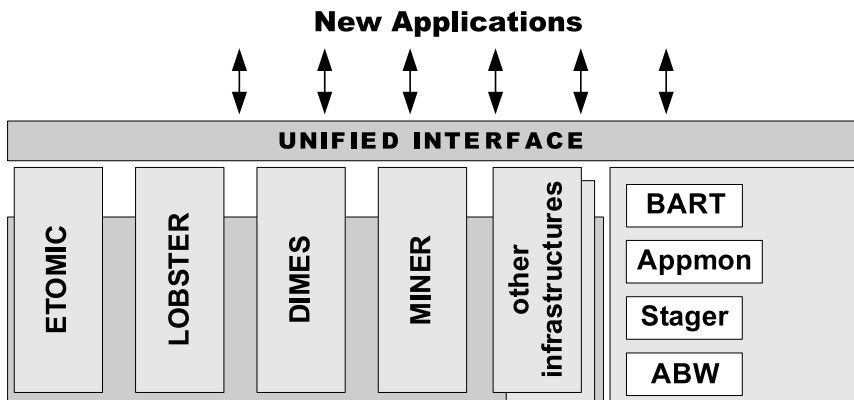


Figure 1: MOMENT Unified Interface.

### Semantics Behind

For unified access to such diverse data sources, it is both a necessary and a challenging task to define a common standard for data formats and semantics. Examples requiring unification are units and names; for example, while one data source stores one-way delay as 'delay' in microseconds, another may store it as 'owd' in nanoseconds. Through the mapping to a common ontology (type hierarchy), infrastructures themselves can be left untouched. For this reason, MOMENT has created a domain ontology for describing measurement data as well as accompanying metadata. The concept of measurement metadata has already been introduced in earlier projects such as DatCat and MOME. The MOMENT consortium defines the metadata structure in the Web ontology language (OWL).

### Service-Oriented Architecture

Due to the diversity of data sources, the MOMENT project uses a service-oriented architecture. This allows the independent development of services, and connection through Web service technologies, fulfilling one of the main

requirements of the project: to provide an open interface for new measurement infrastructures and other future data sources. Single data sources provide an interface of their choice, eg through the Simple Open Access Protocol (SOAP), Representational State Transfer (REST), or even direct SQL or FTP access to the data, which is transformed into services by the MOMENT mediator. A service-oriented approach is also used inside the mediator. By connecting single services to a service chain one can reuse existing operations on various data sources. In this context we want to highlight two core services of the MOMENT mediator, the analysis service (ANAS) and the workflow service (WFS), schematically depicted in Figure 2.

The analysis service allows different kinds of data analysis, ranging from simple histograms to complex calculations such as correlating data from different sources, discovering possible intersections and providing options towards more complete results derived from the broader coverage of the global Internet that MOMENT provides. This




Figure 2: MOMENT Simplified Architecture.

correlation may result in possible new metrics of interest for the future Internet, deriving from the completeness of both the measurement points and the measurement data.

Another feature of the MOMENT mediator is the workflow service. Due to the volatility of the Internet, measurement data are usually only valid for a limited time, so automatic mechanisms including data collection from different data sources and multi-level analysis are helpful services for complex data analysis. The workflow service implements a scripting service that allows workflows to be predefined, stored and repeatedly executed on the different data sources and services provided by the mediator. Workflows are represented in the business process execution language (BPEL).

### Handling Large Data Amounts

Internet measurements often involve large amounts of data, which may cause problems for busy servers or tiny clients. The MOMENT mediator uses continuations in order to divide data streams into manageable pieces, along with a handshake protocol between client and server so that very large data sets can be conveniently handled.

### Links:

- <http://www.fp7-moment.eu>
- <http://miner.salzburgresearch.at>
- <http://www.barttool.org>
- <http://www.ist-lobster.org/>
- [http://lobster.ics.forth.gr/~appmon/public\\_sensors](http://lobster.ics.forth.gr/~appmon/public_sensors)
- <http://www.netdimes.org>
- <http://www.etomic.org>

### Please contact:

Felix Strohmeier  
Salzburg Research  
Forschungsgesellschaft mbH/AARIT,  
Austria  
Tel: +43 662 2288 443  
E-Mail:  
[felix.strohmeier@salzburgresearch.at](mailto:felix.strohmeier@salzburgresearch.at)

Martin Nilsson  
Swedish Institute of Computer Science  
(SICS), Sweden  
Tel: +46 8 633 1574  
E-mail: [martin.nilsson@sics.se](mailto:martin.nilsson@sics.se)

Demetres Antoniadis  
ICS-FORTH  
Tel: +30 2810 391654  
E-mail: [danton@ics.forth.gr](mailto:danton@ics.forth.gr)

# Privacy-Aware Network Monitoring

by Peter Dorfinger, Carsten Schmoll and Felix Strohmeier

*Collecting information as required for network operation also gathers personal information from users sending their data over the network. Existing network-monitoring applications do not take user privacy into consideration by design. We present a framework that allows these applications to operate in a privacy-preserving environment.*

Network monitoring is a central component behind a stable operational network. It is used to guarantee the security of the network infrastructure or to validate service-level agreements. To extract the necessary information, network operators collect data sent by individual users with monitoring applications. Often the privacy of the users is not seriously taken into account, and the gathered information may contain a significant amount of personal content (Figure 1). Since IP addresses constitute private information, each captured packet contains private information. If the payload is also captured then passwords (email, ftp etc), e-mail content or VoIP calls can simply be regenerated from the captured packets. Since the employees of the operator, at least those working in network operation, will generally have access to these traces, the privacy of the users is threatened. As user privacy legislation extends further into the Internet, network monitoring on the future Internet will only be able to take place if privacy can be guaranteed. Since network monitoring is necessary, the PRISM (PRIVacy-aware Secure Monitoring) project consortium has developed a framework that allows network monitoring to take place in a manner that ensures user privacy in the future Internet.

This solution is based on a two-tiered approach as shown in Figure 2. In the front-end block the traffic is captured. The front end processes the traffic with the goal that no privacy information is handed over to the back end. Depending on the purpose of the monitoring application, individual fields of the packets may be, for example, deleted, randomized, summed up or anonymized. The information is sent to the back end where it is stored. Per request or for on-the-fly monitoring without intermediate storing, the information from the front end is transformed to the input format of the external monitoring application and exported to it. Fields of the input data that are not received from the front end



Figure 1: Collected traffic of a privacy unaware network monitoring application.

can be filled with random values. The application operates on a privacy-preserving modified/reduced data set strictly tailored to its needs.

Transforming a legacy monitoring application into a privacy-preserving one will result in different front-end processing for each individual monitoring purpose. One such example, the

Skype traffic detection engine of TSTAT (Transfer Control Protocol Statistical and Analysis Tool), is described here in detail. TSTAT currently captures all the traffic on a link, meaning any information transported on the link can be rebuilt from the packets, including information relevant to privacy. Further IP addresses are pri-

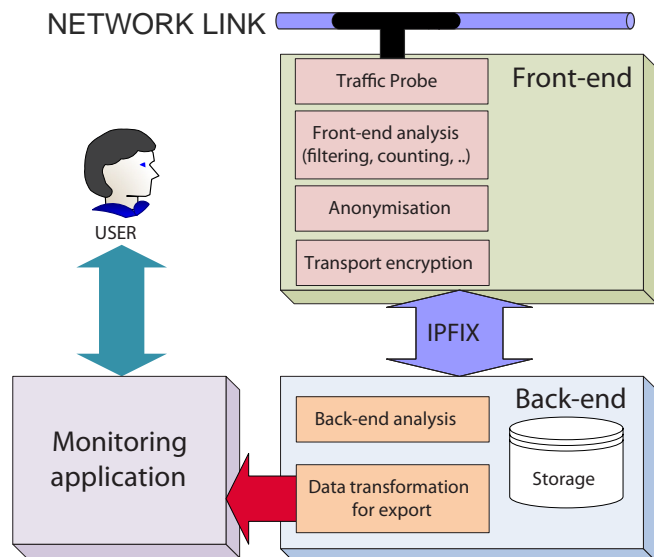


Figure 2: PRISM architecture.

vacy-sensible data, especially since they indicate who communicates with whom. Skype uses payload-encrypted packets for transportation. Thus the first task in the front end after capturing the packets will be to filter for packets where the payload is encrypted. Encrypted payload can be handed over to an external application because the information in it cannot be rebuilt. All other payload must be randomized or stripped because, depending on the protocol, it may contain private information in plain text.

One important aspect is the question of how to handle the common 5-tuple attributes (src/dst IP address, protocol and src/dst port number). Since the IP addresses are the most privacy-sensitive of these fields, they should be removed or at least remapped in a non-reversible way. It is recommended that src and dst IP addresses be mapped together; this means that each observed IP address is mapped to a different new value for

each different src/dst IP address pair. In this way the mapping space is much larger and single addresses cannot be reverse-mapped easily by injection attacks. The protocol number is generally non-critical and can be kept. Port numbers could be removed (set to zero) in our example (with the exception of the well-known Skype port number 12340 for outgoing Skype traffic), effectively hiding all other port numbers since they are not needed for this analysis.

The information is then sent to the back end and stored. On a request the back end generates a tracefile and hands it over to TSTAT. TSTAT can now perform Skype traffic detection on a tracefile that contains all Skype traffic and is free of sensitive information.

Future work in the PRISM project will be the implementation of the proposed framework and the adaptation of existing monitoring applications. The focus

of the work for involved ERCIM members lies in the adaptation of monitoring applications and on traffic anonymization.

#### Links:

<http://www.fp7-prism.eu>  
<http://tstat.tlc.polito.it/skype.shtml>

#### Please contact:

Peter Dorfinger  
Salzburg Research  
Forschungsgesellschaft mbH/AARIT,  
Austria  
Tel: +43 662 2288 452  
E-mail:  
[peter.dorfinger@salzburgresearch.at](mailto:peter.dorfinger@salzburgresearch.at)

Carsten Schmoll  
Fraunhofer Institute for Open  
Communication Systems - FOKUS,  
Germany  
Tel: +49 30 3463 7136  
E-mail:  
[carsten.schmoll@fokus.fraunhofer.de](mailto:carsten.schmoll@fokus.fraunhofer.de)

## Breaking the Weakest Link: Becoming a Trusted Authority on the Internet

by Marc Stevens

*When weaknesses are found in cryptographic protocols or algorithms on which the everyday security of the Internet relies, it is important that they are replaced by more secure alternatives. This is clearly emphasized by the recent case of MD5, an algorithm commonly used to create digital signatures in which severe weaknesses were found in 2004. The continued use of MD5 by several leading Certification Authorities (CAs) on the Internet enabled our team to become a rogue CA itself, triggering quick and adequate responses from the affected CAs and major Internet browsers.*


When you visit a Web site whose URL starts with 'https', a small padlock symbol appears in the browser window. This indicates that the Web site's digital identity is verified using a digital certificate issued by one of a few trusted CAs. To ensure that the digital certificate is legitimate, the browser verifies its signature using standard cryptographic algorithms. It is one of these algorithms, known as MD5, which can be misused. MD5 is vulnerable to an attack called 'chosen-prefix collisions', which essentially allows one to manipulate any two files so that they would receive the same MD5-based signature. In principle this allows fraud by copying a signature between two such 'colliding' files, for instance by copying a CA's signature from an issued certificate to a rogue cer-

tificate. Such an attack becomes a real threat when the rogue certificate is in fact a CA certificate instead of a regular end-user certificate.

MD5 is an algorithm that splits a file or certificate into small fixed-size blocks that are iteratively processed to update the small internal state that MD5 maintains. A chosen-prefix collision attack manipulates a few sequential blocks in both files to remove all differences between the two internal states in an iterative process. The two internal states will thereafter be equal and remain equal when the remaining blocks are also equal for both files and in the end lead to the same digital signature. The number of blocks available for these manipulations is severely limited in the

case of certificates as issued by real CAs, greatly increasing the computational cost normally required for a chosen-prefix collision attack.

Our new mathematical improvements allow more differences per message block to be removed in a very flexible manner and allow new trade-offs between the computational and memory costs of various parts of the attack. These improvements can greatly reduce the computational costs for short chosen-prefix collisions depending on the available memory. Combined with the use of a cluster of 215 PlayStation3s, performing like 8600 PC cores, this enhanced attack can now be performed on certificates in only a single day whereas previously it would take this cluster two



*Figure 1: The differences between our two crafted digital certificates in the computation of their signatures. This clearly shows first random differences, then the three blocks with very structured differences of our attack. After this there are no longer any differences, leading to both certificates having the same signature. Picture: CWI.*

months. When making use of terabytes of hard-drive space, this enhanced attack is even possible in a single day on twenty PS3s or on the publicly available Amazon EC2 computing service at an estimated cost of \$2,000.


For the attack to succeed, the validity period and serial number of the issued certificate had to be predicted correctly, as both will be generated by the signing CA. We found one particular CA for which this could easily be done and after a few attempts we succeeded. The correctly predicted issued certificate thereby contains a signature that is also valid for the colliding rogue CA certificate, promoting it into a CA certificate trusted by all major Internet browsers. With the example rogue CA we managed to demonstrate that a critical part of the Internet's infrastructure was not safe. In combination with known weaknesses in the DNS (Domain Name System) protocol, this could have opened the door for virtually undetectable phishing attacks. For example, without being aware of it, users could be redirected to malicious sites that appear exactly the same as the trusted banking or e-commerce websites they believe they are visiting. The Web browser could then receive a forged certificate

that will be erroneously trusted, and users' passwords and other private data can fall into the wrong hands.

To prevent any damage from occurring, the certificate we created had a validity of only one month – August 2004 – which expired more than four years ago. The only objective of our research was to stimulate better Internet security with adequate protocols that provide the necessary security. All affected CAs responded quickly and migrated to more secure alternatives to MD5 only hours after our presentation of the rogue CA at the 25C3 security conference in Berlin on December 30, 2008. The leading CAs together with the major Internet browsers are pushing stronger security measures for all CAs, since this CA infrastructure is only as strong as its weakest link.

More research is needed into alternatives to MD5 on possible weaknesses and the real security they provide. The National Institute of Standards and Technology (NIST) is holding an open competition to select a new standard SHA-3 to replace MD5 and its current alternatives.

The team consisted of Marc Stevens (Cryptology Group, CWI), Alexander Sotirov, Jacob Appelbaum (Noisebridge, The Tor Project), Arjen Lenstra (EPFL), David Molnar (UC Berkeley), Dag Arne Osvik (EPFL) and Benne de Weger (TU/e).



*Three members of the team that created the rogue CA, in front of the PS cluster: Marc Stevens, Arjen Lenstra, and Benne de Weger. Picture by Joppe Bos.*

#### Links:

<http://www.cwi.nl/en/node/839>

<http://www.win.tue.nl/hashclash/rogue-ca/>

#### Please contact:

Marc Stevens, Cryptology Group, CWI, The Netherlands

Tel: +31 20 592 4103

E-mail: [Marc.Stevens@cwi.nl](mailto:Marc.Stevens@cwi.nl)

## ResiliNets: Resilient and Survivable Networks

by David Hutchison and James P.G. Sterbenz

*The ResiliNets initiative is an umbrella for a number of projects in resilient future Internet architecture. It aims to understand and improve the resilience and survivability of computer networks, including the Internet.*

Society increasingly relies on computer networks in general, and on the Internet in particular. Consumers rely on networks for access to information and

services, personal finance and for communication with others. The Internet has become indispensable to the routine operation of businesses and to

the global economy. The military depends on network-centric operations and warfare. Governments depend on networks for their daily operation,

service delivery and response to natural disaster and terrorist attacks.


The consequences to disruption of the network are therefore increasingly severe, and threaten the lives of individuals, the financial health of businesses, and the economic stability and security of nations and the world. With the increasing importance of the Internet, so follows its attractiveness as a target for attackers, whether they be recreational or professional hackers, terrorists or those intent on information warfare.

Lancaster University in the UK and the University of Kansas in the US are

Science Foundation (NSF) programmes FIND (Future Internet Design) and GENI (Global Environments for Network Innovation).

The ResumeNet project (Resilience and Survivability for Future Networking: Framework, Mechanisms, and Experimental Evaluation) hosts original research work that aims to systematically embed resilience into the future Internet. The project work has three main dimensions.

First, at framework level, work will elaborate on the  $D^2R^2+DR$  framework and quantify the impact of different types of challenges and failures on nor-



*Figure 1: ResiliNets first employs a two-phase strategy  $D^2R^2+DR$ . The first phase  $D^2R^2$  begins with defence, making the network as resistant as possible to challenges. Inevitably however, a network will be threatened and it must be able to detect this automatically. It will then remediate any damage to minimize the overall impact, and finally will recover as it repairs itself and transitions back to normal operation. The second longer-term phase DR consists of diagnosing any design flaws that permitted the defences to be penetrated, followed by a refinement of network behaviour to increase its future resilience. From this strategy, we derive a set of design principles leading to resilient networks.*

developing new network architectures that regard resilience and survivability as critical to the future of our network infrastructure. The ResiliNets initiative aims to understand and improve the resilience and survivability of computer networks, including the global Internet, PSTN (public switched telephone network), SCADA (supervisory control and data acquisition) networks, mobile ad hoc networks and sensor networks.

ResiliNets is an umbrella for a number of projects in resilient future Internet architecture to which we will apply this strategy and principles. These projects are funded by the EU FP7 FIRE (Future Internet Research & Experimentation) programme, as well as the US National

mal network service operation. Metrics, classes of network resilience, policies and ways to negotiate them are at the epicentre. The outcome of this should improve network engineers' understanding of resilient network services.

Second, the project will look into the mechanisms and algorithms to add to the network infrastructure so that the requirements identified at the framework level can be met. The network infrastructure, both the end-to-end transport infrastructure and the service layer, including servers and data storage equipment, are the subject of work. The focus is on basic functions and technologies that constitute the building blocks of the  $D^2R^2+DR$  framework;

nevertheless, some study cases (scenarios) of network service failure are considered in order to exemplify the synthesis of these blocks into a working system.

Finally, experimentation takes a leading role, both in terms of resource consumption and significance to the project objectives. Specific scenarios focusing on particular network service provision are composed from failure types and resilience mechanisms. The aim is to demonstrate both the applicability of the systematic approach as it unfolds, and the benefits for several aspects of the network functionality. An equally important task will be to provide feedback on the fine tuning of individual mechanisms in the network. The experimentation facilities will come from in-house test beds and facilities developed by the project.

The PoMo (Post-Modern Internet Architecture) project recognizes that the emerging global Internet consists of heterogeneous realms using a variety of technologies (such as optical and wireless) with a number of trust and policy boundaries. In response, we are designing a minimalist Internet working protocol that operates over these realms, and provides strict separation of concerns and inclusion of explicit mechanisms in support of all foreseeable policies that would influence network layer behaviour.

The GpENI (Great Plains Environment for Network Innovation) is deploying flexible and programmable network infrastructure to support interdisciplinary research and development of future Internet architectures, as part of the NSF GENI program. GpENI will provide an experimental platform for a number of research projects, including ResumeNet and PoMo.

**Link:**  
<http://wiki.ittc.ku.edu/resilinet>

**Please contact:**  
David Hutchison  
Lancaster University, UK  
Tel: +44 1524 510331  
E-mail: [dh@comp.lancs.ac.uk](mailto:dh@comp.lancs.ac.uk)

James P.G. Sterbenz  
The University of Kansas, US  
Tel: +1 508 944 3067  
E-mail: [jpgs@ittc.ku.edu](mailto:jpgs@ittc.ku.edu)

# Standardized Testware for Internet-Based Telecommunication Services

by Bostjan Pintar, Axel Rennoch, Peter Schmitting and Stephan Schulz

*Future telecommunication networks will use Internet technology. The European Telecommunication Standardization Institute (ETSI) is developing a standardized test infrastructure to enable industry to provide high-quality open telecommunication services. The testware is based on the Testing and Test Control Notation (TTCN-3) that has been successfully deployed worldwide in a variety of industrial domains.*

Today telecommunications is one of the key areas in which technology is bringing mobility and freedom to daily life. The market is moving in two dimensions from ISDN/PSTN to Internet-based networks, and from previous national monopolies to a variety of network and/or service providers. European residents are looking for economic solutions but still require the reliability and quality of conventional telecom services. As one of the major players, the European Telecommunication Standardization Institute (ETSI) has developed a technical base for future telecommunication services that will allow interworking between existing PSTN-based (ISDN) and new IP-based networks (SIP/IMS).

The main differences between past and future telecommunication networks can be identified by technical issues such as packet-switched data transport, text-based encoding and a mix of public and private network elements. Obviously these innovations need to be checked to ensure functionality and quality constraints. ETSI has therefore initiated the development of a collection of testing


artefacts and events that accompanied the definition of telecommunication standards. This catalogue of ETSI 'products' includes not only the definition of test objectives and test programs with libraries implemented with the Testing and Test Control Notation (TTCN-3), but also interoperability test events.

TTCN-3 is a standardized test definition and implementation language applicable for test modelling and automatic test execution. Several commercial and open-source tools are available to compile TTCN-3 to, for example, Java or C programming code. The tests defined by ETSI using TTCN-3 are freely available from the ETSI websites. One example of ETSI's work in the area of test specification for Internet technology is the development of standardized conformance and interoperability tests for assessing IPv6 core, mobility, security and transitioning functionality. More recently ETSI has started a number of Specialist Task Forces (Figure 1) to develop tests for basic calls (call establishment, release) and supplementary services (call diver-

sion, conferences etc), for both interworking between PSTN- and IP-based networks and between multiple IP-based networks. More concretely, this work focuses on SIP-ISUP interworking, IMS services that fall under national regulatory requirements, and conformance, network integration and interoperability tests for IMS core networks.

Testing experts at ETSI are developing and validating their test programs together with the latest telecom equipment from ETSI members. This helps to assure the quality of the future telecommunication technology at an early time of system development. The expert teams have also been applying new test methods including model-based testing (MBT) approaches.

The efforts at ETSI are supported by its members from industry and academia as well as the European Commission. They enable a base for the open communication between different national and international telecom equipment providers and operators worldwide. The ETSI approach of open standards gives



*Figure 1: STF Support Unit Director Alberto Berrini and Senior Technical Expert Stephan Schulz from ETSI (middle and right in 2nd row) with the STF368/369 experts from Fraunhofer FOKUS, Testing Technologies, Apica IT, FSCOM, Italtel and Sintesis in front of the ETSI headquarter in Sophia Antipolis, France.*



a common technology base in order to facilitate future Internet-based telecommunication. It provides a general and non-proprietary value for Europe since communication will have a significant influence on our future.

#### About ETSI

ETSI produces globally applicable standards for Information and Communications Technologies (ICT), including fixed, mobile, radio, converged, broadcast and Internet technologies, and is officially recognized by the European Union as a European Standards Organization. For more information, please visit: <http://www.etsi.org>

#### Links:

<http://www.ttcn-3.org>  
<http://www.ipt.etsi.org/>  
<http://www.etsi.org/WebSite/OurServices/plugtests/home.aspx>  
<http://portal.etsi.org/stfs/process/home.asp>

#### Please contact:

Bostjan Pintar  
SINTESIO, Slovenia  
Tel: +386 41311993  
E-mail: [pintar@sintesio.org](mailto:pintar@sintesio.org)

Axel Rennoch  
Fraunhofer FOKUS, Germany  
Tel: +49 30 3463 7344  
E-mail: [axel.rennoch@fokus.fraunhofer.de](mailto:axel.rennoch@fokus.fraunhofer.de)

Peter Schmitting  
FSCOM, France  
Tel: +33 608515187  
E-mail: [peter.schmitting@fscom.fr](mailto:peter.schmitting@fscom.fr)

Stephan Schulz  
ETSI, France  
Tel: +33 4 9294-4964  
E mail: [stephan.schulz@etsi.org](mailto:stephan.schulz@etsi.org)

## Network Description Tools and Standards

by Freek Dijkstra, Jeroen van der Ham and Ronald van der Pol

*Setting up high-speed network connections – light paths – is still a manual effort taking two to three weeks. The network engineers need a clear picture of the network topology in order to plan and configure light paths. The Network Markup Language (NML) standardizes network topology and state information. The University of Amsterdam and SARA Computing & Networking Services in the Netherlands are contributing to this effort, with the ultimate goal being to automatically set up and manage light paths.*

In recent years, network engineering has become much more challenging. Most European research networks are offering light paths, dedicated network connections with guaranteed quality of service. With the advent of these new network services, researchers are demanding light paths across the globe that cross many administrative networks. Network engineers must keep track not only of changes at the IP layer, but also of daily changes in the underlying circuits.

What is more, the properties of other networks become relevant for the operation of these circuits. For path finding, it is necessary to have the details of different encodings in order to avoid incompatibilities in a path. Finding the shortest light path is highly nontrivial, and a recent study shows that shortest paths can even contain loops. This counter-intuitive behaviour is caused by incompatibilities between different networks that force the shortest path to include network domains that can convert between different encodings. Exchanging information about encodings and topologies between network domains is essential for path finding.

Monitoring existing connections can be cumbersome, as failures in one domain may propagate through other networks, while it is unclear in which domain the failure originated. Traditionally, network operation centres have used dedicated software to manage their own networks, and relied on phone and e-mail to exchange details about their network with associated networks. This approach no longer works for present-day research networks that offer light paths to their users. Automated exchange of topology and state information between networks is necessary.

#### Network Description Language

The network description language (NDL) is a schema that allows domains to describe and exchange information about their networks. It was developed by the University of Amsterdam within the Dutch GigaPort project. The main selling point of NDL is that while it allows for a precise interoperable description of the network topology, it also permits the individual networks to control the exchange and confidentiality of information.

Networks are traditionally represented as graphs, with nodes that represent either devices or whole domains. NDL deviates from this simple model and provides a much richer schema, which allows the description of devices, interfaces, domains, adaptations and more. The inclusion of adaptations allows explicit descriptions of multilayer networks, including hybrid networks that combine technology such as Ethernet, SONET (Synchronous Optical Network) and WDM (Wavelength Division Multiplexing) in the same infrastructure.

#### Standardization

Research networks are collaborating to make both practical and formal agreements on the exchange of data. The Global Lambda Integrated Facility (GLIF) provides an informal platform to agree on procedures, while the Open Grid Forum (OGF) provides a more formal platform to standardize data formats and protocols for information exchanges.

Developers of NDL contribute to the Network Markup Language workgroup (NML-WG) that standardizes a data

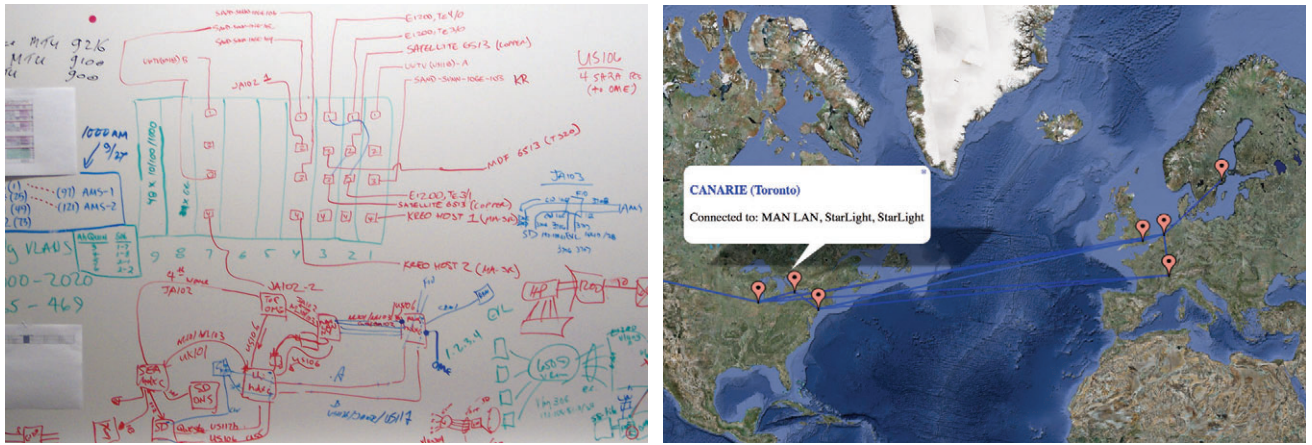


Figure 1: Currently, engineers use whiteboards. With NDL, topologies will be automatically generated. Source: J. van der Ham.

format for topology description. Other contributors include developers of the perfSONAR monitoring software, developers of the common Network Information System (cNIS), in use at the European research network GÉANT2, and many others.

The University of Amsterdam (UvA) and the Dutch computing centre SARA have developed preliminary tools based on NDL. Some of these tools are shown in the figures and are in use for the Dutch research network SURFnet 6, as well as by international network exchange facilities such as NetherLight in Amsterdam.

In SURFnet 6, network engineers must deal with a number of databases to perform their tasks, including an incident

ticket database, fibre database, configuration management database and customer database. Many of the data between the databases overlap and some tools, such as path finding, use information from multiple databases. The goal of NDL is to allow easy integration of the data, while each tool still maintains its own data.

Besides collaborating on a formal standard that incorporates the best of all contributions, SARA and UvA continue to build open-source tools for managing their networks. One of the future goals is to integrate network information with data and computing resources. This should allow grid schedulers and video applications to provide users with optimal performance, taking computing, storage and the network into account.

#### Links:

Global Lambda Integrated Facility (GLIF): <http://www.glif.is/>

Network Description Language (NDL): <http://www.science.uva.nl/research/sne/ndl/>

SARA Network Research Group (NRG): <http://nrg.sara.nl/>

Network Markup Language (NML): <http://forge.gridforum.org/sf/projects/nml-wg>

#### Please contact:

Freek Dijkstra  
SARA Computing & Networking Services, Amsterdam, The Netherlands  
Tel: +31 20 5928068  
E-mail: [Freek.Dijkstra@sara.nl](mailto:Freek.Dijkstra@sara.nl)

#### Invited Article

## Delay-Tolerant Bulk Internet Transfers

by Nikolaos Laoutaris and Pablo Rodriguez

*From its conception the Internet has been a communication network, meaning its development has been driven by the assumption that connections and data transfers are sensitive to delay. Spatial optimization in the form of routing has therefore been the main tool for improving services offered by the network. Temporal optimization, in the form of scheduling, has been limited to millisecond-second scales and aligned with the requirements of interactive delay-sensitive traffic. In recent years however, the network has been progressively shifting from communication to content dissemination. Unlike communication, content dissemination can often tolerate much larger delays, eg in the order of hours. This higher tolerance to delay allows scheduling to go beyond congestion avoidance. Here, we briefly illustrate how to use store-and-forward scheduling to perform bulk data transfers that may be impossible or, under current pricing schemes for bandwidth, prohibitively expensive.*

#### Bulk Data

Residential and corporate bulk data have fuelled an unprecedented increase in overall Internet traffic over the last few years. On the end-user side, these

include high-definition movies from commercial Web sites or peer-to-peer (P2P) networks, large-scale software updates and remote backups. Adding to this, data centres hosting cloud comput-

ing applications exchange large amounts of synchronization, accounting and data-mining traffic, while large corporate and government organizations contribute increasing numbers of eco-

nomic, engineering and scientific datasets. Internet service providers (ISPs) and manufacturers of networking equipment are finding it increasingly difficult to keep up with this traffic without substantial new investments that are hard to find in the competitive ISP market. The once euphoric belief of infinite network capacity, triggered largely by the fibre glut of the 90s, has thus been quickly replaced by headlining news on the discriminatory practices of ISPs against P2P traffic.

### Bulk Bottlenecks under Flat-Rate Pricing

Several ISPs throttle P2P traffic from flat-rate residential customers during peak hours in order to free capacity for interactive traffic that is valued more highly by most users. Popular P2P applications like BitTorrent tend to be targeted, but in the future it may be any other bulk application that becomes

shaded area behind the two rates. Notice here that had it not been for the combined effect, each throttling behaviour alone would have allowed for much higher volumes (individual areas under the solid and dashed lines). In the extreme case of sender and receiver pairs with long, non-overlapping valleys, the transfer could be throttled across the day. Such situations can occur with end points in remote time zones, or within the same time zone but on ISPs of different types, eg, a residential access ISP peaking in the evening and a corporate access ISP peaking at noon.

### Bulk Bottlenecks under Percentile Pricing

Similar problems can arise under the 95-percentile pricing often applied to corporate customers or hosting services that pay based on (nearly) peak usage. Such pricing is justified by the fact that the cost of networking equipment

evening peak hours. The end result could be additional transit costs as the bulk flow cannot avoid increasing the peak load.

### Store 'n' Forward to the Rescue

As such bottlenecks (due either to pricing or throttling) become more prevalent, we argue that to restore the performance of bulk transfers and minimize transmission costs will require a new 'Store 'n' Forward' (SnF) service based on 'temporal redirection' techniques. Existing 'spatial redirection' techniques like native and overlay routing perform path selections over short periods to avoid bottlenecks at the Internet's core, but have no way of escaping complex accumulated constraints that can occur in the future via the combination of various bottlenecks in different time zones.

To solve the problem, we propose breaking end-to-end flows into smaller segments and performing SnF scheduling through intermediate storage nodes to achieve the best utilization of the capacity available between two end points for long-lived bulk transfers. Storage nodes decouple the end-point constraints: when a sender is not throttled (charged) then the data is uploaded quickly (cheaply) to a storage node where it accumulates until the receiver can also download it quickly (cheaply). In previous work, we used data from a large transit provider to show that SnF can provide Tbyte-sized daily bulk transfers at low cost, or for free, whereas end-to-end connection transfers, and even parcel delivery services would incur a much higher cost. We have also argued for the benefits of SnF scheduling for residential broadband users. This work provides some initial direction in what we believe to be a promising new field of research on delay-tolerant networks.

### Link:

A full PDF version of this article including references:  
[http://research.tid.es/nikos/images/bulk\\_overview.pdf](http://research.tid.es/nikos/images/bulk_overview.pdf)

### Please contact:

Nikolaos Laoutaris, Pablo Rodriguez  
 Telefonica Research, Barcelona, Spain  
 E-mail: {nikos,pablorr}@tid.es

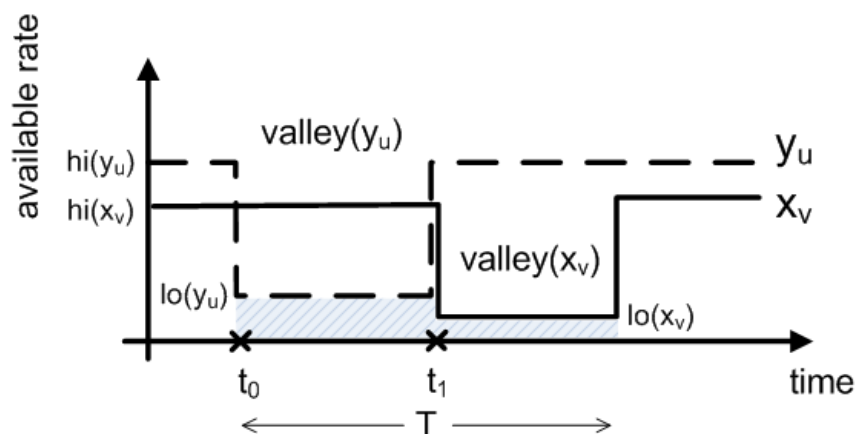


Figure 1: E2E transfers between a sender in ISP  $v$  and receiver in a remote ISP  $u$  having different peak hours can be constantly bottlenecked in  $(t_0, t_0+T)$  due to alternating long-lived bottlenecks in the available uplink rate  $x_v$  and downlink rate  $y_u$  (valley( $x_v$ ) and valley( $y_u$ ) respectively).

popular among flat-rate residential customers. Such throttling bottlenecks introduced by cascades of traffic-shaping devices can have a severe impact on flows across multiple ISPs with different throttling times. In Figure 1 we give an example of a difficult case of combined throttling using hypothetical uplink and downlink rates of a sender and receiver at different access ISPs. In particular, we observe that by chaining such throttling bottlenecks, the combined throttling of receiver (valley( $u$ )) and sender (valley( $v$ )) results in a small transferred volume indicated by the

depends on the maximum load it has to carry with a certain quality of service (QoS). Given that customers pay according to peak traffic, and granted that loads typically exhibit strong diurnal patterns, this leaves much already-paid-for offpeak capacity that can be used to send additional bulk data at no extra cost. Nonetheless, as before, time-zone differences can present a barrier. For example, it might be impossible to use the capacity during the load valley in the early morning hours to send bulk data to a receiver in a distant time zone that is currently going through its

# Resource Management for IPTV Distribution

by Henrik Abrahamsson and Per Kreuger

As part of the SICS Center for Networked Systems we investigate new and efficient ways of distributing television over the Internet Protocol.

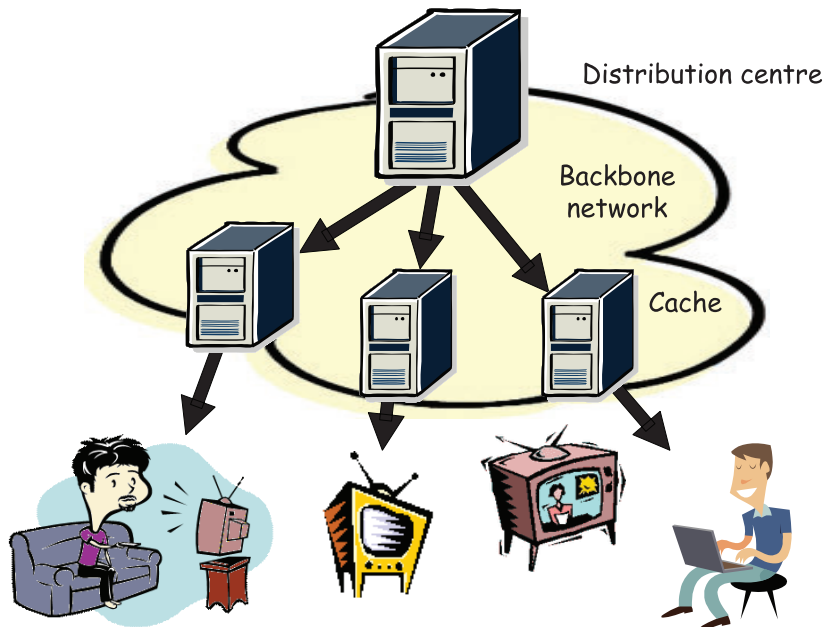


Figure 1: If we regard the TV programs as individual information objects, reruns could be stored in caches closer to the customers, saving considerable amounts of bandwidth.

IPTV, where television is distributed over the Internet Protocol (IP) in a single operator network, has now become reality. Several large telephone companies have become TV providers including AT&T U-verse, Verizon Communications, Deutsche Telekom, TeliaSonera, British Telecom and China Telecom. It is estimated that there are about fifty IPTV operators today with more than 10 million subscribers, and that these figures will increase significantly in the near future.

The TV services provided often include tens of TV channels, sometimes high-definition TV, and video on demand (VoD). While most operators distribute the TV channels over their backbone network as continuous streams, they also often employ a system of peripheral storage devices (caches) located closer to the customers, for time-shift TV and for more permanent storage of popular VoD content.

Most of the programs we see are pre-recorded, with only a small percentage being sent live. Some are recorded the same day, while most are ready for distribution well before they are shown. Many channels also show a lot of

reruns, with prime-time programs being shown again during the night or the next day.

If we regard the TV programs as individual information objects (rather than just part of a streamed channel), the reruns could be stored in the caches closer to the customers. That would save considerable amounts of bandwidth.

A step further would be to pre-distribute recorded content to the caches and release it to the viewers at a later scheduled time. This would allow the network to push the programs out to caches during periods with low network load, typically at night.

With many channels and many hundreds of programs, the issue of when best to transfer the content from a distribution centre to the caches becomes quite a challenging scheduling problem.

In the SICS Center for Networked Systems we have developed scheduling models for IPTV distribution, and verified them against data extracted from real TV schedules. For a case involving sixteen channels over five days we can reduce the total bandwidth requirements

by 21%, just by avoiding sending the reruns over and over again. By scheduling the transfers of content to the caches we can reduce the maximum bandwidth requirement by 25% and perhaps more interestingly, reduce the bandwidth requirement during prime-time periods by 62% – bandwidth that can then be used for video or data traffic.

IPTV also means an evolution of television with new features. For instance, viewers can choose to watch what they want at any time. However, distributing individual TV or video streams to each viewer requires a lot of bandwidth and is a big challenge for TV operators.

Resource management, with techniques for scheduling and pushing content out to caches close to the viewers, is therefore essential for future TV distribution.

#### Link:

SICS Center for Networked Systems:  
<http://www.sics.se/cns>

#### Please contact:

Henrik Abrahamsson  
SICS, Sweden  
Tel: +46 70 7741595  
E-mail: [henrik@sics.se](mailto:henrik@sics.se)

# VISTO: Visual Storyboard for Web Video Browsing, Searching and Indexing

by Marco Pellegrini

*Imagine that searching for videos by specifying visual content is as easy as searching for Web pages by specifying keywords (as is done with search engines like Google and Yahoo). This is the Holy Grail of video searching by content. In the Internet of the Future (IOF), architectural support for multimedia will be influenced by new user-centric search paradigms. The VISTO project is a step towards this goal.*

Video browsing and searching is rapidly becoming a very popular activity on the Web. This has led to the need for a concise searchable, browsable and indexable video content representation, as well as new IOF architectures to support it.

YouTube, the most famous Web-based repository of (short) videos, is growing at breakneck pace. In March 2008, the total number of videos uploaded was 78.3 million, with over 150,000 new videos being uploaded every day. On average each video runs for 2 minutes 46 seconds, meaning it would take about 412 years to view all of the material uploaded on YouTube. Clearly any practical system devoted to indexing video data sets of YouTube's size must function on a time scale orders of magnitude smaller than the duration of the video being processed.

New paradigms of user behaviour induced by new searching tools will have a strong impact on the future of the Internet. Video content is already one of the largest contributors to Internet traffic and several architectural solutions are

being proposed to handle it (centralized server-based, grid-based, cloud computing, peer-to-peer etc). Indexing, searching-by-content and delivery of multimedia data represent a key application driving architectural trade-off choices for the IOF.

Managing large video repositories involves finding solutions for several aspects:

- compression, coding, transmission
- synthesis and recognition
- storage and retrieval
- access (interface, matching user and machine)
- searching (based on machine intelligence)
- browsing (based on human intelligence).

The Visto project is focused on introducing innovative solutions for video indexing to improve the speed and accuracy of searching and browsing. VISTO is intended as a distributed application to be optimized on IOF architecture.

Currently for browsing purposes, static video summarization techniques can be

used. Unfortunately, state-of-the-art techniques require significant processing time, meaning all such summaries are currently produced in a centralized manner, off-line and in advance, without any user customization. With an increasing number of videos in video collections and with a large population of heterogeneous users, this is an obvious limit. In the VISTO project, we employ summarization techniques that produce high-quality customized, indexable and searchable on-the-fly video storyboards. The basic mechanism uses innovative fast clustering algorithms that select, on the fly, the most representative frames using adaptive feature vector spaces. Such techniques are naturally embeddable into distributed scenarios.

Browsing activity: browsing involves an aspect of serendipity in which the user needs to quickly sift through large numbers of videos (say ten or twenty) in order to detect those of interest. Summaries must be made quickly (on the fly) and be of high quality. Users can select basic parameters such as the storyboard length and processing time, as well as advanced ones such as the simi-

*Figure 1: Storyboard of fifteen frames extracted from a fifteen-minute clip of a popular Italian talk-show. The quality of the storyboard is high: each character appearing in the clip has been captured at least once, and in visually diverse poses. The fifteen frames were selected from the 30,000 frames in the clip, in 20 seconds, starting from raw HSV representation, on a 3.5 GHz CPU machine.*



larity metric, the dynamic (motion-flow) content and other features. Moreover, use of personalized criteria is made possible by the fact that very little pre-computation is needed. Preliminary objective and subjective evaluations of basic principles show that storyboards produced on the fly can be of high quality.

Searching activity: the user may encounter a frame/scene during video-browsing that is of particular interest (eg a city skyline) and be curious as to whether other videos contain a similar scene/frame (eg a different take of the same city skyline). At this point the user could extract new and unexpected knowledge from a comparison of different videos having this key common scene/frame.

Personalization of browsing and searching: given a video that is of interest for a user, browsing involves being able to quickly assess the most interesting (relevant, typical or unusual) frames/scenes in the video in order to decide whether it is worth watching in its entirety (or which portions are worth viewing). Users should be allowed to select at view-time basic parameters such as storyboard length and waiting time. How-

ever more advanced summarization criteria need to be supported. For example users must be allowed to specify the dynamic (motion-flow) content they are interested in (eg a moving car is different from a parked one). Further, the underlying similarity metric used for selecting the representative frames should be biased using implicit or explicit user requirements.

Videos can be considered to be sequences of still images (with sound), and techniques for handling large collections of still images (pictures) might be the first line of attack. However, we do not consider this approach to be suitable. A single video corresponds to thousands/millions of frames (depending on its duration) that, with the exception of scene changes, are locally highly similar to one another. Thus the sheer volume of data poses scalability problems to approaches based on indexing single frames (or a dense blind sample of them) as single still images. Instead, in the VISTO project, we intend to develop video representations and indexing techniques that take dynamic components of the video data as prominent in the representation. The aim is to simultaneously exploit the spatial and

temporal coherence of video data, not only to attain compression, but also to make searching fast and efficient. Attaining both goals simultaneously is challenging.

Future activities will involve setting up a complete network-aware software/hardware/conceptual architecture able to cope with high throughput demands in a P2P style of computation, taking advantage of IOF architectural designs. Awareness of network capabilities, which change over time, is particularly important for use on, for example, wireless hand-held devices.

The VISTO project started in 2007 at IIT-CNR as a collaboration between researchers of the Institute for Informatics and Telematics, CNR, Pisa, the University of Modena and Reggio Emilia, and the University of Piemonte Orientale.

**Link:**  
<http://visto.iit.cnr.it/>

**Please contact:**  
Marco Pellegrini  
IIT – CNR, Italy  
E-mail: [marco.pellegrini@iit.cnr.it](mailto:marco.pellegrini@iit.cnr.it)

## The New Role of Humans in the Future Internet

by Daniel Schall and Schahram Dustdar

*In most cases, service-oriented architecture is realized using Web services technology. At the Vienna University of Technology, we have implemented a platform enabling humans to provide services. We foresee important applications that will be based on Human-Provided Services and software services.*

Service-oriented Architectures (SOA) are typically comprised of software services. Many collaboration and composition scenarios involve interaction between human actors as well as software services. Current tools and platforms offer limited support for human interactions in SOA: we therefore present Human-Provided Services (HPS) and the HPS framework. In particular, the aim of the HPS framework is to:

- offer a service registry maintaining information related to human and software services
- enhance service-related information by describing human characteristics and capabilities

- define interaction patterns using Web services technology so that human actors can efficiently deal with interactions.

HPSs are offered by human actors. Web services technology is used to describe HPSs and to enable interaction with real people. The advantage of HPS is that these services can be used in different compositions and Web-based collaborations.

A possible scenario of human and software services is shown in Figure 1a. A human can define an activity, eg ‘review document activity’, which is transformed into a Web service inter-

face. Such interfaces are typically described using the Web Services Description Language (WSDL). The same standards can be used to describe HPSs and software (Web) services.

Via a graphical user interface - a Web 2.0 portal hosted by the HPS framework - the end user can create HPSs without having to understand XML or Web services technology.

The role of humans in the future Internet and SOA is thus not limited to consuming services (Figure 1b); services can also be provided by human actors. HPS unifies humans and services, because a service can be provided by a

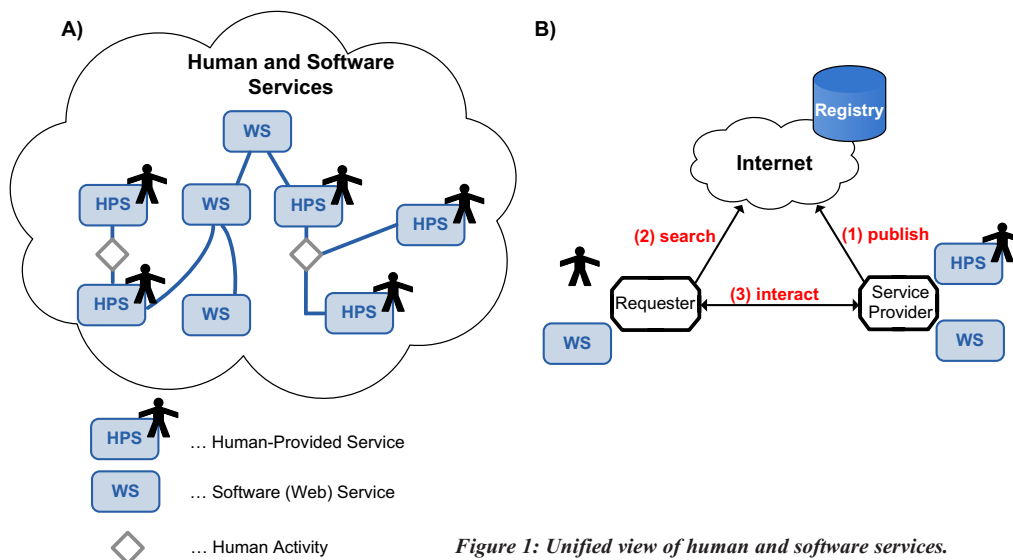


Figure 1: Unified view of human and software services.

human actor or implemented as a software service. Three steps are performed when using HPS:

- 1) Publish: the user can create an HPS and publish the service on the Web using a registry. Publishing a service is as simple as posting a blog entry on the Web. It is the association of the user's profile with an activity depicted as a service.
- 2) Search: the requester can perform a keyword-based search to find HPSs. Ranking is performed to find the most relevant HPS based on, for example, the expertise of the user providing the service.
- 3) Interact: the framework supports automatic user interface generation. HPS can be used in interactions between humans and in interactions between software services and HPSs.

The HPS framework is comprised of three layers (Figure 2). The first layer contains data collections to maintain user profile information, service descriptions (WSDL and related information) and interaction rules. These rules can be created by the user to define event-condition-action patterns such as pre-filtering of interactions, messages etc.

The second layer contains user tools including a graphical user interface for designing services (a simple tool to create HPSs without programming code), the activity management to organize and structure activities, interactions and related messages. The service management allows the user to modify his/her HPSs.

The third layer enables interaction between requesters and HPSs. The lookup is used to find suitable services,

with matching HPSs being ranked and returned (comparable to using search engines on the Web). The access layer is a proxy service dispatching messages that are exchanged in HPS-based interactions.

The middleware services are implemented in Java/J2EE and data collections are managed with XML databases (to manage WSDL- and XML-type definitions) and MySQL to manage interactions and related messages.

The user tools are mainly implemented in C# and ASP.NET with the support of automatic user interface generation using XForm (XML forms) technology. These forms are generated based on XML descriptions (eg WSDL).

The framework offers APIs (lookup, registry), enabling integration with other platforms.

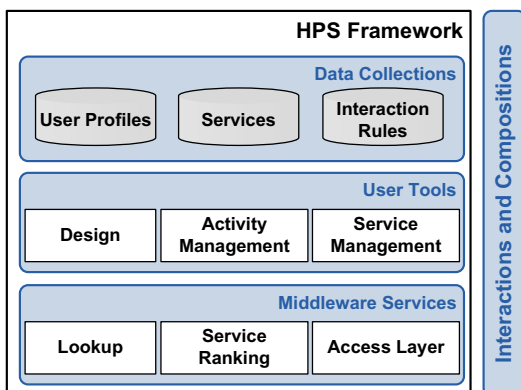


Figure 2: HPS framework.

**Human-Provided Services**

- Document Review/Translation
- J2EE Consultant
- Java Expert
- Lawyer

**Compositions Human and Software Services:**

- Task Force
- Expert Team
- Service/Social Community
- Human Computation

**Link:**

[http://hydra.infosys.tuwien.ac.at/prototyp/HPS/HPS\\_index.html](http://hydra.infosys.tuwien.ac.at/prototyp/HPS/HPS_index.html)

**Please contact:**

Daniel Schall  
Vienna University of Technology / AARIT, Austria  
Tel: +43 1 58801 18453  
E-mail: [schall@infosys.tuwien.ac.at](mailto:schall@infosys.tuwien.ac.at)

Schahram Dustdar  
Vienna University of Technology / AARIT, Austria  
Tel: +43 1 58801 18414  
E-mail: [dustdar@infosys.tuwien.ac.at](mailto:dustdar@infosys.tuwien.ac.at)

# Orchestrating your Surroundings

by Pau Giner, Carlos Cetina, Joan Fons and Vicente Pelechano

*The Internet of Things envisions real-world environments that are augmented with digital services. With ever greater numbers of devices being added to our surroundings, simplicity is greatly appreciated by users. With the purpose of preventing service behaviour from becoming overwhelming, our project is devoted to dynamically organizing services according to user needs.*

Imagine a future in which your fridge announces to you the recipes that can be prepared with the available goods, your TV tells you that your favourite program is beginning, the book you want to start reading is suggesting you try other similar books; and all of this is happening at the same time. Clearly living in such an environment on a daily basis would be annoying. On the other hand, if services behave in a completely auto-

matic manner (without requiring human input), users can feel that their environment is out of their control, which is also undesirable.

Since user attention is a valuable but limited resource, we believe that an environment full of embedded services must behave in a considerate manner, demanding user attention only when it is required. In the same way that a

musical orchestra requires a conductor to indicate who should play and the tempo to be followed, we propose to incorporate orchestration techniques in Internet-of-Things (IoT) environments, in order to achieve a balance between automation and user participation.

Our project consists in a method for building IoT environments that are automatically reorganized in order to offer the ‘degree of obtrusiveness’ required by the user. For example, several devices could be gradually used in order to make the experience of waking up less shocking (see Figure 1). We make use of model-based techniques to face the development of such systems from a technology-independent perspective. The method comprises the following steps:

1. Business process modelling. The activities in which the users are involved and their temporal dependencies are captured by means of business process modelling techniques. In particular, we use Business Process Modelling Notation (BPMN), a graphical notation standardized by the Object Management Group (OMG), since it is well understood by business analysts and provides facilities for deriving executable solutions from the BPMN models.

2. Implicit interaction design. Each task is analysed to determine the level of obtrusiveness it requires. In particular we are focusing on two factors that affect the design of implicit interactions: initiative and attention. Initiative defines whether the service is activated automatically or by the user. The attention factor relates to the user’s awareness of a service, defining whether the service performs in the foreground or the background of user attention. Table 1 illustrates the requirements for the various tasks involved in the example process of waking a hotel guest.

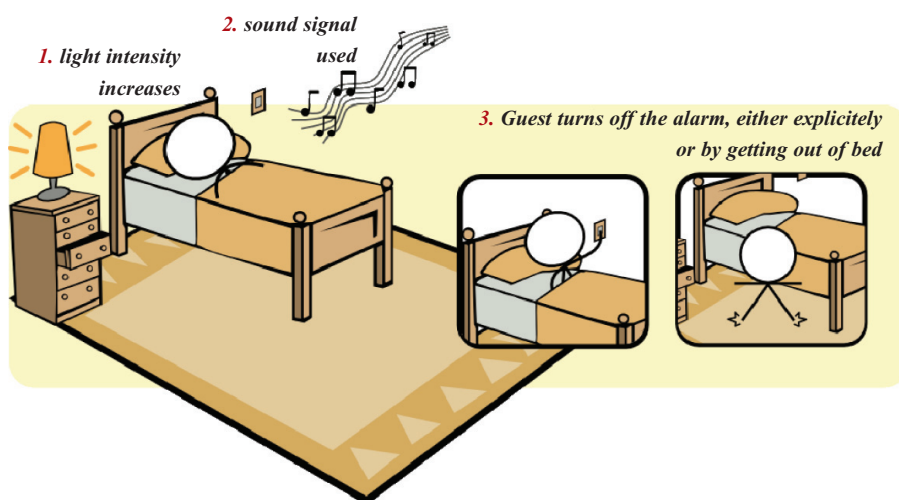


Figure 1: Example scenario illustrating the process of waking a hotel guest.

Task	Obtrusiveness	
	Initiative	Awareness
Wake up	<b>Proactive.</b> The system triggers the activation of the mechanisms for waking the user.	Increasingly from <b>background to foreground</b> . The goal is to make the user aware in a gradual manner (using lights first, then sound).
Turn the alarm off	<b>Proactive.</b> The user can decide to turn off the alarm, either because he/she decides to get up or because he/she consciously decides to continue sleeping.	<b>Foreground.</b> The deactivation of the alarm is something that should be consciously done by the user.
Get out of the bed	<b>Reactive.</b> The user activates the sensors that detect that he/she is getting out of bed. Alarm is turned off as a consequence.	<b>Background.</b> Sensor activation is not intended to be conscious for the user.

Table 1: Analysis of the required obtrusiveness for different tasks.



3. Deployment in an orchestration engine. Business process execution technologies have proven successful when applied to coordinating networked services, especially in Service-Oriented Architectures (SOAs). We have extended these technologies, which have already been applied in the 'traditional' Internet, and applied them to the Internet of Things. In particular, we have developed middleware for integrating automatic identification technologies such as RFID or QR codes with business process execution engines based on WS-BPEL (Web Services Business Process Execution Language).

4. System reconfiguration. A model-based reconfiguration engine is used to automatically adjust the system architecture to the requirements presented as models. To achieve this, the system architecture is based on components and communication channels that can be dynamically established among them for achieving a particular behaviour. In this way, a service can transit from silent behaviour to behaviour that is noticeable by the user as a reaction to the business process evolution.

This work is part of a four-year research project at the ProS research centre,

which is devoted to the development of new methods for the production of software in different areas. The techniques presented here have been successfully applied to scale environments with real technology in our research lab, with the measurement of acceptance by real users being the next step in our research.

**Link:**

<http://www.pros.upv.es>

**Please contact:**

Vicente Pelechano

Universidad Politécnic de

Valencia/SpaRCIM, Spain

E-mail: [pele@pros.upv.es](mailto:pele@pros.upv.es)

## An Autonomic Home Networking Infrastructure

by Thomas Luckenbach, Mario Schuster and Marc-Oliver Pahl

*Next-generation home networking environments will contain a variety of Internet-ready devices or embedded systems, which will result in increased complexity for the end user. New methods are therefore required to build autonomic networking infrastructures that enable auto-configuration and self-management of the networked elements and keep technical details hidden from the user.*

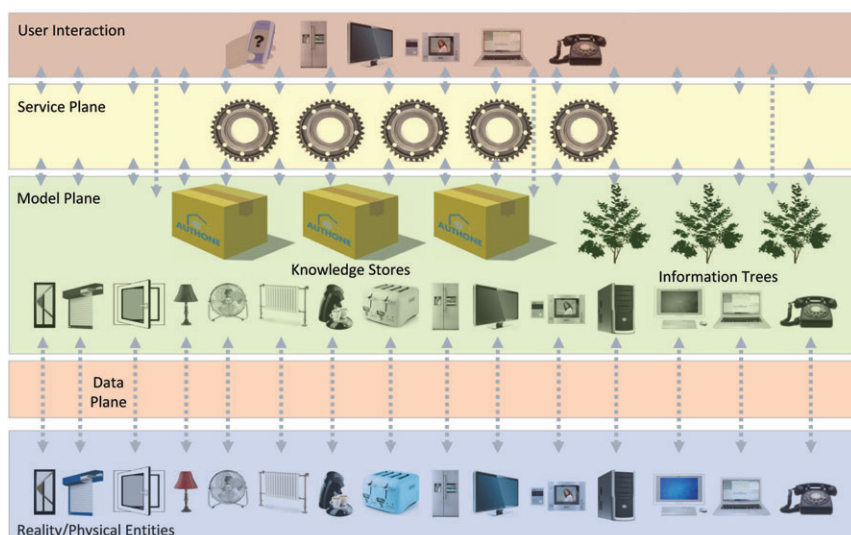
In future Internet-based home environments a variety of devices such as PCs, smart phones, networked appliances and embedded systems will be organized into infrastructures and will act together to form a new type of service provisioning platform. Today's typical home environments often contain a few of these Internet-ready devices; in the future there will be even more, including embedded systems like sensors and actuators that are also accessible via

Internet technology (eg TCP/IP, HTTP, Web services etc). While on the one hand this gives users the freedom to access a rich variety of devices in the vicinity of their homes, on the other it increases the depth of knowledge required to understand what is happening in such a complex networking environment. Novel methods for autonomous networking as well as self-configuration, self-management, self-organization, self-protection and self-

healing (self-x) of the networking elements can take responsibility for observing and controlling the home network without directly involving end users.

For this purpose, the German national project Autonomous Home Networking (AutHoNe) has been started in 2007. The main goal of the project is to enable autonomous networking technologies in home and industrial environments. The project takes into account a variety of currently available or future technologies that can operate in those environments. Another aspect of AutHoNe is to provide users with a consistent and transparent view and approach to accessing information and services in the home. This includes the development of mechanisms for the intuitive definition of rules, constraints and other parameters that influence the autonomous behaviour of the networking elements.

AutHoNe introduces a hierarchical concept, whereby low-level devices and communication protocols are integrated into a high-level scheme that is capable of executing so-called knowledge services to observe the entire network, and a collaborative decision-making process



**Figure 1: Hierarchical three-tiered architecture for modelling and using knowledge in a home network environment.**

based on the gathered knowledge of the network. This also includes redundant behaviour of the services as well as the ability to distribute rule and profile descriptions within the network.

For this purpose, the different information and data structures coming from the low-level networking elements - which are often completely different and proprietary - must be adapted to a unified knowledge model to ensure that the autonomic methods of the knowledge services can operate in a generic and technology-independent way. AutHoNe therefore defines a three-tiered approach (see Figure 1) that includes:

- a data plane representing the basic information coming from the basic network elements at the bottom layer, as well as communication not addressing the AutHoNe framework (eg non-control information like video streams etc)
- a model plane representing the high-level abstraction of the information as mediator layer
- a service plane as top layer enabling the autonomous control of the network based on the knowledge abstraction.

The knowledge model is defined through state-of-the-art description techniques (eg XML) containing basic information of the network elements as well as rule and profile descriptions and

meta-information. An important requirement is that the knowledge model is available to all nodes in the network that run knowledge services in a distributed and redundant way. All the AutHoNe services operate on the modelled world. Since services have models too, they can be puzzled together in many different ways allowing service providers to mash up their services from other services. The selection of adequate protocols for the communication and model and knowledge representation is still in progress. Our model-centric design could be the key for a future home network rich in useful services that can be provided as easily as iPhone applications are today.

Another important issue is the adaptation and integration of the information coming from the low-level networking elements to the knowledge model. For instance, IP-based devices can provide information about their capabilities and their state through typical network management protocols (eg Simple Network Management Protocol - SNMP). Embedded systems like sensors or actuators can provide information through other methods (eg ZigBee profiles, SensorML/TransducerML descriptions, or IEEE1451 data sheets). Ongoing work in the AutHoNe project will consider these approaches and implement suitable adaptation mechanisms in order to support current and future technology

embedded in an overall knowledge-based home-networking infrastructure.

The project AutHoNe is being partly funded by German Federal Ministry of Education and Research under FOKUS grant agreement no. 01BN0702. The project is being carried out as part of a CELTIC cluster with the grand agreement no. CP3-001 within the EUREKA framework.

#### Links:

<http://www.authone.de>  
<http://www.celtic-initiative.org/Projects/AUTHONE/>  
<http://www.fokus.fraunhofer.de>  
<http://www.net.in.tum.de>

#### Please contact:

Thomas Luckenbach  
 Fraunhofer FOKUS, Germany  
 Tel: +49 30 3463 7245  
 E-mail:  
[thomas.luckenbach@fokus.fraunhofer.de](mailto:thomas.luckenbach@fokus.fraunhofer.de)

Mario Schuster  
 Fraunhofer FOKUS, Germany  
 Tel: +49 30 3463 7302  
 E-mail:  
[mario.schuster@fokus.fraunhofer.de](mailto:mario.schuster@fokus.fraunhofer.de)

Marc-Oliver Pahl  
 Technische Universität München, Germany  
 Tel: +49 89 289 18004  
 E-mail: [pahl@net.in.tum.de](mailto:pahl@net.in.tum.de)

## Load-Balancing Energy Usage of Household Appliances

by Lennart E. Fahlén

*The combination of ambient/ubiquitous interface technology and Internet-based energy load-balancing tools with a consumer focus, holds promise for increased energy savings and control.*

### Informed consumers save energy

Incentives to be energy-aware are increasing on many levels: economical, social, political and environmental. We know that consumers want to save energy if they can, and studies show savings between 5-15% can be achieved simply by providing direct feedback. However, when it comes to managing energy usage in the household, today's consumers are pretty much left to find their own solutions, such as reading electricity meters and

using timers to switch equipment on and off. To address this, the new concept of the 'Internet of Things' provides promising ways to integrate the control of energy into everyday life without the need for drastic lifestyle changes. Computer technology makes it possible to transform complex information into a more accessible form, resulting in improved consumer energy awareness and control and assisting consumers in reducing energy usage.

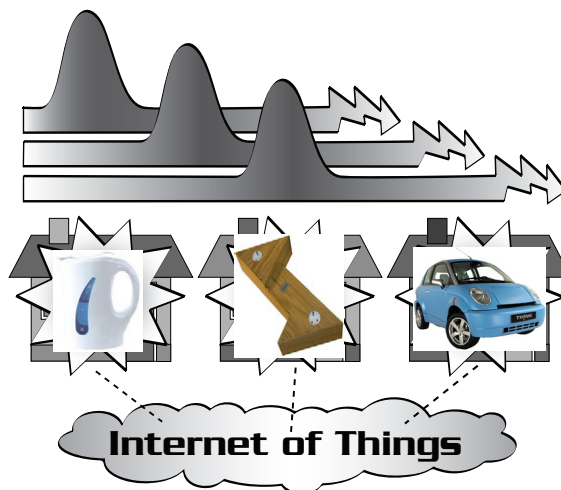
### The load-balancing mechanism

Shifting electricity usage (by time and space) is a desirable goal for many energy-saving initiatives. A further goal is to explore the possibilities inherent in aggregating energy usage into cooperating pools for load-balancing purposes. Pooling usage and reducing peak loads allows an aggregation of use with the potential, for instance, to negotiate lower average rates (possibly offset by higher peak load rates at the margin).

Consider this scenario. A number of appliances within a physical area (homes, apartments, parking lots etc) are connected together. By having an agreement not to exceed a determined maximal energy usage at different times throughout the day, the peak levels and the cost of energy can be optimized. The charging units (or rather the users mediated through the units) negotiate in order to find a suitable timing slot in which to start charging. If for some reason the allocated slot is unsuitable, it is possible to enter a state of manual negotiation or an auction. An appliance in this context could be a variety of things, ranging from various kitchen utensils, through heating devices, to electrical vehicle chargers.

A simple load-balancing system for pooling usage across several units/users could work in the following way. First, a micro-controller controls the hardware modules inside the appliance (force feedback module, current-measuring module, power switch, ID mechanism, GPS etc) and communicates with a charging booking service over the Internet. The booking service is connected to a range of charging appliances, and monitors their current and predicted future power usage. The booking service in turn reports to a pooling service, which supervises the activity over many charging units. The pooling service is thus able to construct an aggregated 'profile' of the near-future power usage of connected units. This profile is then continuously communicated back to the booking

*Figure 1: Shifting electricity usage is a desirable energy-saving goal. Another goal is to explore the possibilities of pooling energy usage for load-balancing purposes.*



services, and finally the individual units.

The profile may be biased by any changes in the power suppliers' cost/allowance levels and/or by maximum power usage limits, which may be stipulated in contracts between the power supplier and the subscribers/users of the pooled units.

To the user, the load-balancing mechanism appears as an ambient and physical interface that fits unobtrusively into the everyday surroundings. Without requiring an understanding of complex visualizations and convoluted interactions, the user is made aware of the current and predicted state of the energy consumption in comparison with other subscribers. Furthermore, the user is able to control and negotiate this consumption together with fellow consumers in a socially aware context.

#### Tea kettle and power strip

Prototypes and demonstrators of the concepts outlined above have already been built and successfully demonstrated. One prototype consisted of tea kettles, where users were able via force feedback means to schedule the boiling of water. Another prototype took the form of a 'power strip', where control and scheduling was guided by tactile and positional means, and by the use of sound, vibration and visual indications. For low-level control and distribution, the PART middleware system was used.

This work was sponsored by The Swedish Energy Agency's Energy IT and Design program.

**Please contact:**  
Lennart E. Fahlén  
SICS, Sweden  
Tel: +46 8 633 15 39  
E-mail: lef@sics.se

## The Internet of Vehicles or the Second Generation of Telematic Services

by Markus Miche and Thomas Michael Bohnert

*The commercial introduction of broadband wireless communication technology such as UMTS or WiMAX facilitates the connection of vehicles with the Internet and hence with a broad range of service providers. It thus paves the way for the second generation of telematic services, which is referred to as Vehicle-to-Business communication. It allows for the enhancement of existing enterprise applications as well as the identification and realization of novel business models, but requires an appropriate infrastructure on top of the aforementioned communication technology. This article presents an integration architecture with which to realize an efficient and secure information exchange between vehicles and back-end services, and points out the potential of the Internet of Vehicles.*

Based on comprehensive research activities and the support of the European Commission, Intelligent Transport Sys-

tems (ITS) will soon reach market introduction. Vehicle-to-Vehicle (V2V) communication centres particularly on

application scenarios in the area of traffic safety and efficiency, as for instance Green-Light Optimal Speed Advisory

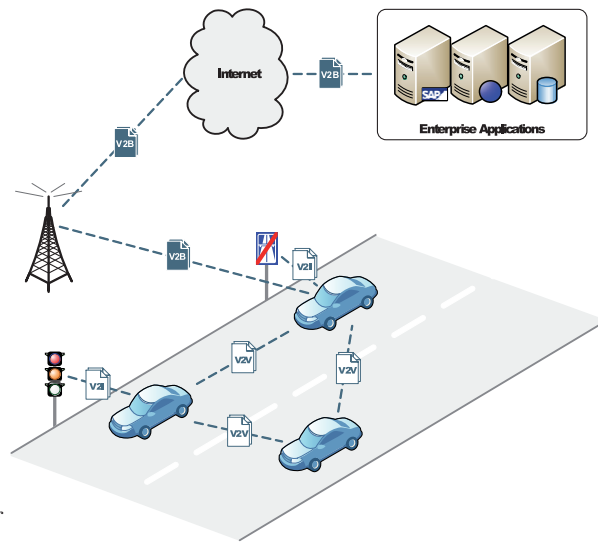


Figure 1: Overview of vehicular communication.

or Traffic Jam Ahead Warning. The recently defined standard IEEE 802.11p for inter-vehicular communication, designed according to the specific requirements of V2V interaction, constitutes an essential step towards this next phase.

However, since message exchange between vehicles is based on vehicular ad hoc networks (VANETs), V2V interaction is subject to large network effects. Due to the short reach of inter-vehicular connections, a certain market penetration is required to provide value for drivers.

Vehicle-to-Infrastructure (V2I) communication represents the second field of ITS. It covers the information exchange between vehicles and roadside infrastructure equipped with wireless communication technology such as traffic lights or warning signs for road works. V2I facilitates the interaction of vehicles and roadside units to enhance the aforementioned application scenarios. Moreover, those units may be used as additional hops to augment the reach and thus the overall value of inter-vehicular communication.

Besides V2V and V2I, the commercial introduction of UMTS, WiMAX and similar broadband wireless access technology, as well as the proliferation of WLAN-based (IEEE 802.11) multihop access networks and hotspots, paves the way for the second generation of telematic services. Such technology allows vehicles to be connected to the Internet and therefore with a broad range of service providers. In addition to infotainment and entertainment services, this Internet of Vehicles facilitates the interconnection of vehicles with com-

mercial business services provided by enterprise applications. For that reason, the second generation of telematic services is referred to as Vehicle-to-Business (V2B) communication. V2B covers a multitude of application scenarios, as for instance remote management of embedded systems, Vehicle Relationship Management (VRM), and the novel insurance concept of usage-dependent tariff models such as Pay As You Drive. Figure 1 illustrates a scenario that includes the three depicted

vehicular communication types V2V, V2I, and V2B.

The main challenges for an efficient V2B communication are the uncertain connectivity of vehicles, the frequency of changes in network topology due to the vehicles' mobility, as well as their resource-constrained onboard units. To account for these demands, SAP Research is currently working on a service-based, distributed integration architecture based on state-of-the-art technology to realize a flexible, scalable and secure V2B interaction covering a multitude of vehicles and diverse (legacy) enterprise applications.

Based on the complementary architecture patterns Service-Oriented Architecture (SOA) and Event-Driven Architecture (EDA), the V2B integration architecture specifies two dedicated components: the Vehicle Integration Platform (VIP) as a back-end system and the Back-end Integration Manager (BIM) as the related in-vehicle component (Figure 2).

The VIP is a Web services-based back-end system that covers multiple components to realize an efficient information exchange between vehicles and enter-

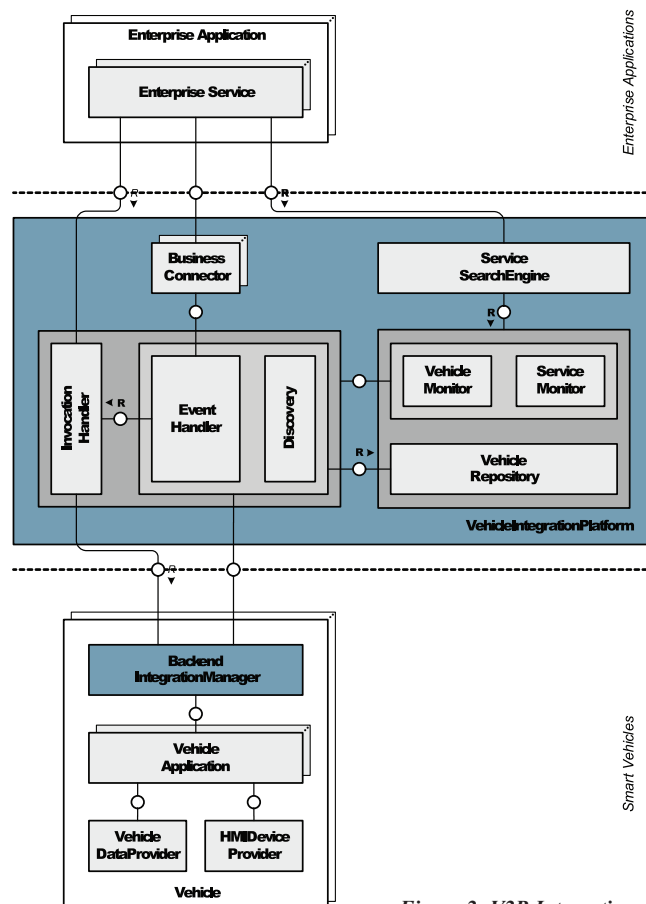


Figure 2: V2B Integration Architecture.

prise applications. It comprises a message broker (EventHandler) that enables a topic-based publish/subscribe interaction style, thus enhancing the decoupling of vehicles and enterprise applications. Moreover, the Invocation-Handler component facilitates the invocation of vehicle applications and provides a buffering mechanism to account for their intermittent accessibility. Finally, the VIP embodies several components to discover, organize (eg identity management) and search for functionality offered by vehicles.

The in-vehicle component BIM is designed to connect a vehicle with the VIP and hence with enterprise applications. It encapsulates vehicle applications as Web services and provides additional functionality such as message caching, message prioritization and the rescheduling of message delivery. The BIM thus applies the Devices Profile for Web Services (DPWS) specification, which allows for secure Web services operations on resource-constrained devices.

V2B communication provides invaluable potential for both service providers such as vehicle manufacturers and insurance companies, and the drivers as end customers. Thus, for example, VRM covers the periodic polling and analysing of in-vehicle systems and sensors by a vehicle manufacturer. Such a large amount of product life-cycle data might be used to improve the product design and quality or to enhance the planning of production and procurement. Hence, VRM might result in a reduced forecast error and a decrease in safety stocks. Moreover, information about driving behaviour can be analysed to provide personalized product and service offerings, which comprise warranty and maintenance as well as value-added services. An effectively operated VRM might further result in an intensified and enhanced Customer Relationship Management (CRM) due to the potential growth in customer loyalty.

The ability to make use of a comprehensive set of vehicle data during the entire life cycle as well as the offering of cus-

tomers-specific services facilitates not only enhancements of existing services but also the identification and realization of novel business models. The aforementioned usage-dependent insurance concept Pay As You Drive represents only one example of the overall potential. Due to its modular design and the application of state-of-the-art technology, our V2B integration architecture provides a valuable foundation for further developments to the Internet of Vehicles.

**Links:**

<http://www.pre-drive-c2x.eu>

<http://www.comesafety.org/>

**Please contact:**

Markus Miche

SAP Research CEC Zurich

Tel: +41 58 871 7833

E-mail: [markus.miche@sap.com](mailto:markus.miche@sap.com)

Thomas Michael Bohnert

SAP Research CEC Zurich

Tel: +41 58 871 7801

E-mail:

[thomas.michael.bohnert@sap.com](mailto:thomas.michael.bohnert@sap.com)

## Web-Enabled Tracking Operations in Distributed Supply Chains

by Zsolt Kemény and Elisabeth Ilie-Zudor

*A substantial number of the challenges to be tackled by the future Internet are posed by collaborative communities such as supply chains or production networks, where massive amounts of product-related information are shared in a selective and secure manner. One of the key application areas in this domain is tracking and tracing, ie, keeping track of production-related entities of interest and their interaction with identifiable objects in their environment. The international project TraSer – Identity-based Tracking and Web Services for SMEs – addresses research and development efforts in this area.*

Aside from case studies, tutorials and instruction materials, the main output of the project is an entry-level tracking and tracing platform. This has been designed especially for small-scale users in heterogeneous environments such as constantly changing production communities with frequent occasional collaboration actions. As things stand, this user range lies outside the scope of most off-the-shelf solution kits for tracking and tracing of products and product data.

TraSer supports tracking and tracing on the level of individual items (as opposed to the account-based view of purely quantitative stock levels), and relies on Web services for communication. Item-

related data are maintained by TraSer servers, which communicate with each other if queries or updates are forwarded or broken down to distributed components. TraSer clients serve as interfaces to 'the rest of the world', including human operators, automated checkpoints or other components of the given IT infrastructure (Figure 1).

Aside from its focus on small enterprises, a number of features distinguish the TraSer solution from most other comparable frameworks. First, TraSer uses an internal identifier notation (ID@URI), where each unique ID already contains the address of the access point for the related services,

thus eliminating the need for a comprehensive lookup service and allowing a decentralized allocation of identifiers. (Note that TraSer can also adapt to unique identifiers in existing numbering schemes by mapping them onto its internal notation in a variety of ways.) The TraSer platform is essentially independent of the physical ID carrier (eg barcode, RFID) and can also be run with purely electronic identifiers for data tracking.

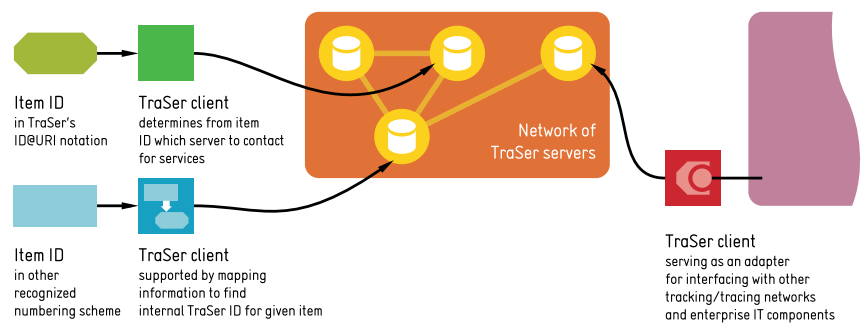
TraSer is also free and open-source, and its implementation relies on a Java-based open-source Web service framework. Publicly available interface specifications will allow users to easily

implement their own special building blocks, such as clients tailored to specific purposes, or to reimplement a TraSer solution in a development environment of their choice. This feature is of special importance for small enterprises that act as suppliers in multiple production networks and must conform to multiple communication and track-and-trace specifications.

TraSer offers useful functionalities for handling item-centric data such as advanced search and aggregation support, views of historical data, and back-tracking of information. These facilitate special actions such as focused recall campaigns or the location of certain goods in the production chain. At the moment, such actions are still seriously hampered in heterogeneous production networks due to the lack of transparency and commonly accepted interfaces.

TraSer allows a flexible definition of data models and facilitates this with its XQuery-based interfaces, a feature which is of special importance for enterprises of high product variability (as opposed to the majority of today's track-and-trace applications which are optimized for large production volumes and less diversity). Furthermore, users can easily extend or overlay existing product data maintained by other manufacturers, which adds further flexibility and freedom in data model planning.

In addition, a special release of the solution platform was adapted to tracking product data files: not common in commercially available tracking platforms.



**Figure 1:** Users gain access to item-related services in a TraSer network by various methods of unique identification and with specialized clients using different interfaces. The freedom to add customized clients means a TraSer-based solution can be adapted and coupled with other components and systems.

With completion planned for June 2009, and approaching the end of its three-year life span, the TraSer project has now advanced to the stage of pilot implementations representing various targeted user groups. These were selected in order to provide useful practical feedback about industrial needs that were, in part, unknown even to the users in question. Performed at various research institutions, small enterprises and logistics service providers, the key emphasis of implementation pilots dealt with following cases:

- tracking of product data in design, prototyping and procurement processes
- tracking of product delivery in collaboration with small enterprises, using a minimum of IT resources
- closed-circuit tracking of reusable equipment in research facilities
- closed-circuit tracking of the transportation assets of a country-wide logistics service provider and offer-

ing goods-tracking services for customers.

Upon termination of the project, an industry-ripe release of the solution platform will be available for free download, along with accompanying material such as tutorials, instructions, case studies and educational material for prospective users entirely new to the concept of tracking and tracing. Further information can be obtained at the official project Web site or by contacting the project representatives.

**Link:**

<http://www.traser-project.eu/>

**Please contact:**

Elisabeth Ilie-Zudor; Zsolt Kemény  
SZTAKI, Hungary  
Tel: +36 1 279 6195  
E-mail: [ilie@sztaki.hu](mailto:ilie@sztaki.hu),  
[kemeny@sztaki.hu](mailto:kemeny@sztaki.hu)

## Knowledge-Based Collaboration Patterns in Future Internet Enterprise Systems

by Gregoris Mentzas and Keith Popplewell

**A challenge for enterprise systems in the future Internet of services is to address the sharing of knowledge in virtual organizations. SYNERGY has developed a service-oriented solution for knowledge-based collaboration in such virtual organizations.**

The last decades show a clear trend in business: away from big, comprehensive trusts which can cover all stages of a value creation chain, and away from long-standing, well-established and stable supply chains. Instead, companies

are increasingly focusing on their core business competencies and often enter into flexible alliances for value creation and production. This growing demand for flexibly interactive and efficiently integrated businesses and services has

already led to a significant amount of scientific and technological research on enterprise interoperability. This has already achieved promising results, and has been partly responsible for initial commercial products and service offer-

ings and operational, deployed applications. Nevertheless, these remain at the level of data interoperability and information exchange; they hardly reach the level of knowledge integration, and certainly fall short of knowledge-based collaboration.

We claim that research on collaboration patterns in the event-driven future Internet has the potential to facilitate the satisfaction of these needs and provide the underlying technological infrastructure for supporting adaptive enterprise collaboration through knowledge services. We work on the development of such knowledge services within the SYNERGY project, which aims to enhance support for the successful and timely creation of, and participation in, collaborative virtual organizations by providing an infrastructure and services to discover, capture, deliver and apply knowledge relevant to collaboration creation and operation.

The vision of the SYNERGY project is related to a Web-based service-oriented software infrastructure that will help all kinds of companies that need to engage in collaborative ventures to discover, capture, deliver and apply knowledge relevant to collaboration creation and operation. This will allow them to effectively and efficiently participate in Virtual Organizations (VOs).

The SYNERGY architecture takes up and refines the challenge of the Interoperability Service Utility (ISU) as defined by the Enterprise Interoperability Research Roadmap produced by the European Commission. This challenge is to produce an open, service-oriented platform that allows companies to use an independently offered, intelligent infrastructure support to help in the planning, setting-up and running of complex knowledge-based collaboration.

The business models supported by SYNERGY refer to virtual value networks of enterprises. Through its lifetime, each network will develop project-specific knowledge. This is in part knowledge specific to the network's product or service, and to the processes and technology involved, but it is also related to the current state of the network in its life cycle. In most cases, such knowledge needs to be maintained only for the network and its partners

because it is of no use, and is possibly even very confusing, outside that context. Nevertheless, there may be a need to analyse such knowledge and its evolution to provide improved patterns for the future, thus forming a basis for organizational learning. Such collaboration patterns may then enter the public domain to support future collaborations across, say, an industrial sector. However, there will be cases where it is precisely this knowledge which represents a competitive advantage to the network or partners concerned. These cases need to be identified and services must deliver new learning to a specified and appropriate audience of future users – perhaps only partners in the network generating this new knowledge.

The maintenance of a library of appropriate collaboration patterns, available as process and service templates to be specialized as necessary and applied to network enterprises either as they form or as they subsequently evolve, is central to the support of knowledge-based partner collaboration in a VO. The SYNERGY system will therefore embody structures and services to define collaboration pattern templates and to select (according to the nature of a developing or existing/changing network), specialize and apply such templates.

In SYNERGY we have adopted and extended work on software patterns in the knowledge-based collaboration realm. By adopting patterns as the organizational model of collaboration, users will work in a more complete context for their actions and be burdened by fewer manual integration tasks. At the same time, by dividing collaborative work into distinct activities, users can focus more readily on a particular activity and deal more easily with interruptions by suspending and resuming activities as needed. Collaboration patterns augment rather than replace collaboration services and tools. Through reference-based integration methods, collaboration patterns introduce new views and organizational schemes that cut across service and tool boundaries, thereby increasing the integrity of the representation of work and mitigating scatter.

In SYNERGY we assess the appropriate level of pattern granularity (abstraction) and represent collaboration patterns using ontologies (an ontology for collaboration patterns, codenamed COMPANION, is publicly available). This ontology will be used to represent collaboration patterns as a collection of relationships that emerge between people, the resources they use, and the artefacts they work on, as well as for the

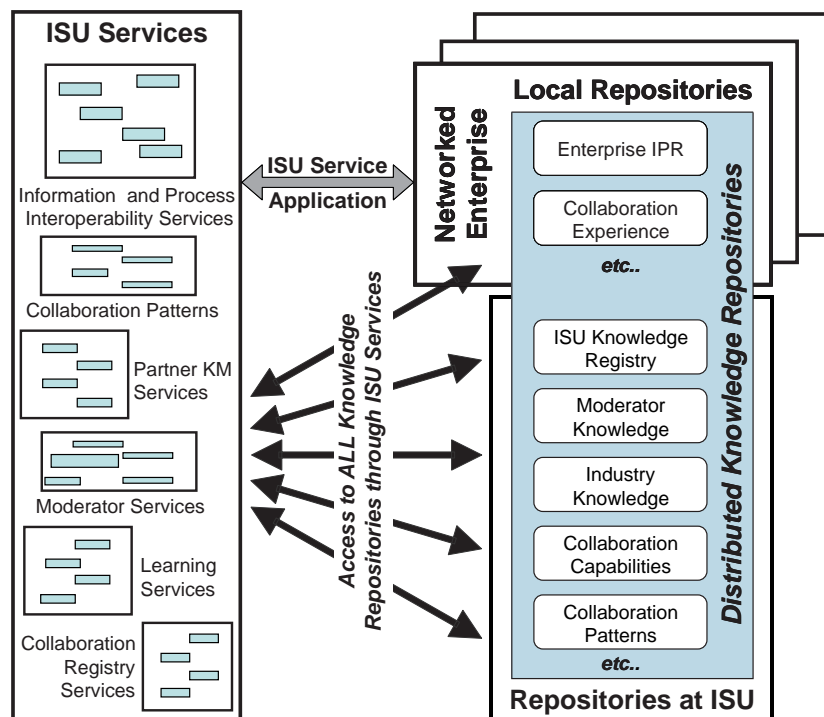


Figure 1: SYNERGY develops Interoperability Service Utility (ISU) services to help networked enterprises discover, capture, deliver and apply knowledge relevant to collaboration creation and operation.

communication, coordination and business processes used to complete their work. We envisage three ways for generating collaboration patterns: manually, from best practice data; semi-automatically, by detecting prominent usage patterns using folksonomy techniques; and by community members themselves, either from scratch or as refinements of existing patterns.

SYNERGY is expected to benefit virtually collaborating organizations in many ways. As a 'learning organization' a collaborating partner is better able to exploit past experience in responding to future opportunities and

in particular to opportunities for participation in collaborative networks. Enhanced sharing of knowledge accelerates and improves network decision making, shortens time to market and reduces network operating costs, whilst improved capture and especially reuse of enterprise and network knowledge through collaboration patterns reduces the cost of repeating work of earlier projects, and of repeating past errors.

#### Links:

SYNERGY: <http://www.synergy-ist.eu/COMPANION>:

<http://www.imu.iccs.gr/ontologies/companion>

Enterprise Interoperability Research Roadmap: [http://cordis.europa.eu/ist/ict-ent-net/ei-roadmap\\_en.htm](http://cordis.europa.eu/ist/ict-ent-net/ei-roadmap_en.htm)

#### Please contact:

Gregoris Mentzas  
Institute of Communication and Computer Systems, National Technical University of Athens, Greece  
Tel: +30 210 772 3895  
E-mail: [gmentzas@mail.ntua.gr](mailto:gmentzas@mail.ntua.gr)

Keith Popplewell  
Coventry University, United Kingdom  
Telephone: +44 24 7688 7017  
E-mail: [cex393@coventry.ac.uk](mailto:cex393@coventry.ac.uk)

## Experiences from the Public Safety Arena for the Future Internet

by Socrates Varakliotis, Peter Kirstein and Steve Hailes

*Telecommunications and information exchange are vital to the operation of emergency response units. The objective of the U-2010 project team is to use existing or future telecommunication infrastructure to provide the most capable means of communication and the most effective access to information to all the parties required to act in case of accident, incident, catastrophe or crisis. We are feeding our experiences with the technology involved in this area into the discussion on the future Internet.*

### Internetworking of Small Devices - The Internet of Things

Emergency response teams are increasingly being equipped with communication devices. To be of use to first responders, these devices must be light, versatile and autonomous, but also miniaturized, capable of sensing the environment, of tracing and tracking resources, and able to convey this wealth of information to civil protection authorities and crisis management centres. To put it simply, Wireless Sensor Networks (WSNs) have already gained momentum in public safety applications. Our approach requires that the devices for public safety be mobile (wearable, portable) and that the communication of their data to remote management teams be done in any possible way ('always available'). The scenario also envisages the use of autonomous subsystems like robots or Unmanned Aerial Vehicles (UAVs) that are IP-enabled.

To enable communication between small devices involved in such missions and to orchestrate effectively the flow of information made available by these devices, we advocate the use of IPv6 as the interconnecting protocol. The multi-

tude of addresses required and its capabilities for node auto-configuration are two of the reasons for the deployment of IPv6 in such scenarios; the improvements in mobility and power usage made in IPv6 protocols over IPv4 for this environment are other important considerations. We also propose the use of light, multipurpose gateways, which interconnect small local clouds of WSNs and temporarily cache measurements for as long as necessary until a communication path to the mission control team is established.

Work done in the area of gateways builds upon some of the developments carried over from the EC-funded RUNES (Reconfigurable Ubiquitous Networked Embedded Systems) project. Here, open-source, small form-factor gateways supporting multiple network technologies such as Ethernet, Wi-Fi and IEEE 802.15.4 are being enhanced with protocol stacks suitable for wireless sensor Internetworking and basic multimedia capabilities. For example, a prototype gateway has been extended by the 6LoWPAN (IPv6 over Low power Wireless Personal Area

Networks) adaptation layer and by the implementation of a lightweight video-streaming service with support for commodity USB cameras.

### Infrastructure: Sensor Network Testbeds

The research approach we adopted is helped by the design, deployment and management of large distributed IPv6-based sensor network testbeds. Much of the work is experimentally trialled on the UCL Heterogeneous Experimental Network (HEN), a large research network deployment including devices ranging in size from low-power sensor nodes to server-grade multi-processor rack systems. We adapt the HEN infrastructure to design and run experiments with large numbers of network nodes, including HEN's customizable sensor testbed comprised of over fifty motes spanning a fairly large lab area with realistic radio channel conditions. We scale such experiments to hundreds of motes and allow integration with external administrative and user planes, such as development offerings by Arch Rock, or entirely proprietary vendors such as Cisco, and other open-source



operating systems, routers and WSN applications.

A number of other projects use HEN for some form of validation, such as work on novel MAC algorithms, for which the ease of creating dense networks with sensor nodes is invaluable.

### WSN Mobility

The mobility requirements of first-response teams are studied in the context of network mobility. Most generic wireless sensor application scenarios assume the sensor nodes are static (wired or wireless). In novel applications, such as emergency response teams using communications equipment and sensor networks, the nodes need be nomadic, ie, nodes that move along with a response team. In networking terms, this translates to mobility paradigms such as NEMO (Network Mobility). Furthermore, in these studies more complex network routes emerge that are best approached by MANEMO solutions (Mobile Ad-Hoc Network Mobility). We are currently investigating both these forms of mobility in WSNs with external partners. We are also investigating what is required in nomadic nodes like robots and UAVs.

In this context, Transport Layer Multipath is currently being investigated as a possible future research area. Enabling a transport layer protocol to create a number of sockets (with distinct paths)




Figure 1: Schematic of remote fireman access in a mobile emergency environment.

and hide them behind what appears to applications as a single socket allows a number of novel interactions. These include raising throughput by load balancing between the sockets, improving reliability through the redundancy provided by multiple paths, and providing a mechanism for mobility by adding and removing paths. In this application, the normal concern of scalability with small numbers of points of failure is less important than in the general case.

Multipath to the sensors could be deployed through multiple paths to multiple gateways to the WSNs, for example. In the context of emergency services, the additional reliability gained by using multiple available backhaul connections (satellite and 3G)

could be essential, while having simple mobility without the requirement for a home agent may also be advantageous. In addition, the potential for the aggregation of bandwidth between different access media could make it possible to offer services requiring high throughput, such as efficiently streaming high-quality video from an emergency site.

### Links:

<http://u2010.eu>  
<http://www.ist-runes.org>  
<http://hen.cs.ucl.ac.uk>

### Please contact:

Peter Kirstein  
 UCL Computer Science, UK  
 Tel: +44 20 7679.7286  
 E-mail: P.Kirstein@cs.ucl.ac.uk

## With Joint Forces - Establishing Situation Awareness for the Future Internet

by Tanja Zseby and Thomas Hirsch

*Several future Internet solutions introduce decision-cycles in the network for protection, management and application support. They require the establishment of situation awareness in network nodes as the basis for making decisions. Fraunhofer FOKUS has developed a Node Collaboration System (NCS) that provides situation awareness based on the collaboration of network nodes. Due to its cross-layer design it also allows the collaboration of services with network nodes to include service feedback to network decision functions.*

As increasing numbers of businesses and communities come to rely on network connectivity, the Internet has mutated behind the scenes into a complex assembly of patchwork solutions. This makes it vulnerable, difficult to manage and hard to adjust to future demands. Future Inter-

net solutions will address network protection, network management and application support that can substitute for this uncontrollable agglomerate of protocols.

Many solutions, like autonomic communication, require decision-cycles in net-

work nodes so that solutions can be applied before problems affect end systems and users. Situation awareness is the prerequisite for making good decisions. Classical decision-cycles describe four phases in which situation awareness is established: the 'observe' and 'orient'

phases, followed by the ‘decision’ phase and finally the ‘act’ phase, in which the decided actions are executed.

Within the project 4WARD, Fraunhofer FOKUS is investigating solutions by which situation awareness can be established as the basis for network decisions. Collaboration of network nodes and between different layers is the key

Distributed Context-Aware Firewall (D-CAF), based on the NCS, that takes valuation reports from services into account when making the decision inside the network about the blocking of packets.

The objective of D-CAF is not to identify attacks, but rather to make an intelligent blocking decision if an overload

ensuring that it is always the least important flows (from the server’s viewpoint) that are blocked. In order to restore normal operation after an overload, the filters are slowly removed when the total traffic reaches a sufficiently low level.

The D-CAF approach lends itself to a distribution of the three functional blocks (reporting; aggregation of reports; firewall configuration) across the network. In non-hierarchical network architectures (Mesh, MANET, SOA), these blocks may be replicated wherever components are able to evaluate users, process data and protect resources. To verify this approach, we plan to deploy and investigate the D-CAF scenario on a larger scale within PlanetLab Europe, a federated large-scale testbed that is extended in the project OneLab. This allows us to set up and accurately control various testing and verification cases. We plan to investigate the various effects of correlating valuation reports from large clusters of protected services, and in particular the timing and delay effects occurring in a realistic simulation.

A further application field for the NCS is the assessment of different future Internet solutions that are based on network decision-cycles. In such a scenario the NCS can be used to simulate different awareness levels with artificial information sources and help to assess key performance indicators such as the costs of collaboration, decision quality etc. Methods for the assessment of autonomic communication solutions are currently standardized within the Autonomic Communication Forum (ACF).

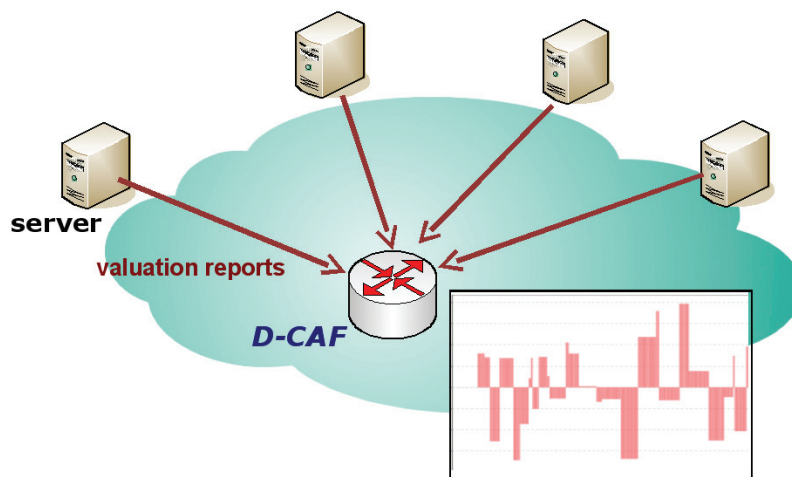


Figure 1: The Distributed Context-Aware Firewall (D-CAF).

to enhancing situation awareness by aggregating information from a variety of different sources.

To this purpose, Fraunhofer FOKUS has developed a Node Collaboration System (NCS) that allows network nodes to access information from neighbouring nodes and end systems. The system provides the means by which to locate information and then access it directly from the source. It is also possible to invoke the on-demand generation of additional information by triggering measurements or post-processing functions.

Based on the NCS, Fraunhofer FOKUS is working on solutions for network protection against overload situations that originate from either legitimate traffic like flash crowds or malicious traffic like denial-of-service attacks. While in such cases it is comparatively easy to detect the overload situation with classical measurement methods, it is not trivial to decide which traffic should be blocked. While it is necessary to reduce traffic in order to protect the network and its users, it would be unfortunate if valuable customers were blocked. We have therefore developed a

situation occurs, regardless of whether the origin is legitimate traffic or an attack. During normal operation, D-CAF learns how important a particular traffic flow is by collecting and analysing valuation reports from its connected servers. A Web site may serve as an example: a user visiting the site with low frequency might get a slightly positive rating, a site scraper with high frequency a slightly negative one. A successful login with a user account receives a high positive rating.

These valuations are mapped to a uniform range per service, which we define as  $[-1.0; 1.0] \in \mathbb{R}$ . This corresponds to the usefulness of the user for the service or business objective. Such reports are sent from all servers via the NCS to D-CAF, where they are aggregated, weighted and used as the basis for the blocking decision.

Based on the valuation reports, D-CAF automatically generates blocking rules that represent the joint opinion of all connected servers. If an overload situation occurs, the first flows to be blocked are those with a low aggregated valuation. Filter rules are applied step-by-step until the load situation normalizes,

#### Links:

Fraunhofer Institute FOKUS:  
<http://www.fokus.fraunhofer.de>

4WARD Project:  
<http://www.4ward-project.eu>

OneLab2 Project:  
<http://www.onelab.eu>

#### Please contact:

Tanja Zseby  
Fraunhofer FOKUS, Germany  
Tel: +49 30 3463 7153  
E-mail:  
[tanja.zseby@fokus.fraunhofer.de](mailto:tanja.zseby@fokus.fraunhofer.de)

# Social Networking for a Pervasive Future Internet: the SOCIALNETS Project

by Stuart M. Allen, Marco Conti, Andrea Passarella and Roger M. Whitaker

*Wireless and mobile devices such as phones, MP3 players, sensors and PDAs are increasingly capable of creating and sharing content. Exploiting their owners' social networks for communication between devices provides a unique way of translating qualitative human behaviour into adaptation for pervasive networking systems.*

The future Internet will be a 'flat' pervasive network in which users will generate and share content dynamically on an unprecedented scale by exploiting the rich capabilities of their mobile devices. At the same time, the social behaviour of the users will be a key driver of their content generation and consumption patterns. Mobile devices will form spontaneous networks that extend traditional (mobile) networks, enabling users to generate and share content anytime and anywhere.

This scenario is targeted by the EU SOCIALNETS project (Social Networking for Pervasive Adaptation), which commenced in February 2008 under the FP7 FET 'Pervasive Adaptation' initiative. It takes a human-centric approach to communication by establishing a social network of relations between fixed and mobile devices. This is defined by human users and their particular actions and behaviour with respect to each other and the technology in the environment. The goal of the project is to establish new paradigms for pervasive networking in the future Internet, where models of the user's social networks are a building block for the design of self-organized communication and content-management services.

To date, emerging methodologies for self-organized networking draw inspiration from biology and the autonomic nervous system. While these approaches can be highly effective, they are not naturally aligned with human behaviour and thus fail to capture the details of whom we interact with, why, and how. In response to this limitation, we are developing a new people-centred approach by exploiting the core characteristics of human behaviour. Social networks are intimately connected with the human and are the basis for the SOCIALNETS paradigm. The project seeks to understand, model and exploit these structures using an interdisciplinary approach involving social anthro-

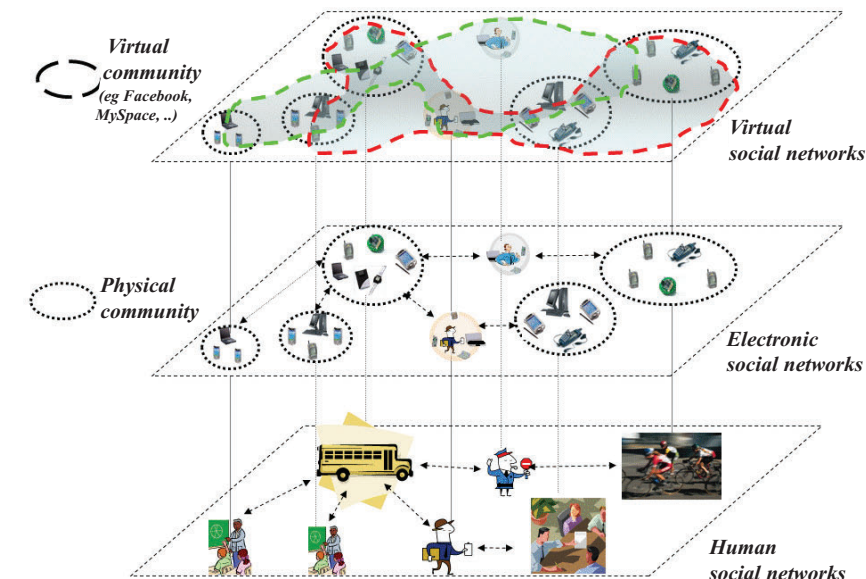


Figure 1: The SOCIALNETS vision.

pology, complex systems, network engineering and computer science.

Within the project we exploit social anthropology research which describes the inner structure and evolution of an individual's social relationships. This evolves from fundamental results describing the structure of social acquaintances between people. In particular, M. Dunbar's number dictates that the typical person can only maintain about 150 acquaintances. Additionally, members of an individual's social circle do not all have relationships of equal closeness. Rather, the group is structured into a series of circles of decreasing acquaintanceship, the size of which scales with a remarkably consistent ratio of almost exactly three. Thus there appear to be important constraints on both the structure of natural social networks and on the quality of relationships. Physics approaches that embrace complex network analysis allow us to derive compact models of these social properties, which can then be used as context information for pervasive networking systems. For example, recent

results achieved in the project allow us to relate the frequency of interactions between users to the level of acquaintance between them, according to simple yet expressive mathematical relations.

In SOCIALNETS, the networking research departs from the convention that fully connected, end-to-end paths should be maintained to enable communication. This is unrealistic for many small and low-powered wireless devices. Reality requires radical rethinking, with network architectures opportunistically harnessing contact between users' devices to exchange multimedia and messages. As stable topological views of the network cannot be maintained in dynamic pervasive networks, SOCIALNETS complements traditional network topology information with context information derived from the social acquaintances analysis. This will allow us to build an Electronic Social Network encompassing users' mobile devices. Conventional communication decisions (eg which node to use as a message forwarder) are taken opportunistically, by dynamically evalu-

ating the fitness of encountered nodes to fulfil the required task (eg how likely an encountered node is to meet the message's eventual destination). Social relationships between people thus translate into social relationships between their devices, with associated measures of trust that are of primary importance in communication. This allows us to derive a social electronic topological view of the pervasive network, which we see as the equivalent, in the future Internet, of conventional network topologies.

On top of this, SOCIALNETS explores novel ways of managing and supporting user-generated content. Rather than a means to establish communication between devices, the network is seen as a shared space for users to store, discover and retrieve content of interest. To enable this vision, we are building on the social electronic topological view described above, to blend it with information about the structure of the users' social communities and their interests. This forms a knowledge base on which content can be replicated, placed and searched for, according to

the expected user's social behaviour and patterns of interactions.

The innovative aspect of the SOCIALNETS project is that instead of exploiting limited features of social networks to improve protocol performance, it harnesses social networks to provide a completely new paradigm. In particular, this will emerge as the ability to self-organize and provide:

- an intelligent, trusted and dynamic knowledge base that can be exploited to coordinate local interactions and effort for data forwarding
- coordination and structured content of relevance to users in massively information-rich pervasive scenarios where there is high diversity
- adaptive trust and security that are community-based and that exploit groups of users with commonality between them.

The seven European partners participating in the SOCIALNETS project are: Cardiff University (Coordinator, UK), National Research Council – CNR (Italy), University of Cambridge (UK),

National and Kapodistrian University of Athens (Greece), University of Oxford (UK), University of Aveiro (Portugal), and Institut Eurécom (France).

**Link:**

<http://www.social-nets.eu>

**Please contact:**

Stuart M. Allen  
Cardiff University, UK  
Tel: +44 29 20876070  
E-mail: [Stuart.M.Allen@cs.cardiff.ac.uk](mailto:Stuart.M.Allen@cs.cardiff.ac.uk)

Marco Conti  
IIT-CNR, Italy  
Tel: +39 050 315 3062  
E-mail: [marco.conti@iit.cnr.it](mailto:marco.conti@iit.cnr.it)

Andrea Passarella  
IIT-CNR, Italy  
Tel: +39 050 315 3269  
E-mail: [andrea.passarella@iit.cnr.it](mailto:andrea.passarella@iit.cnr.it)

Roger M. Whitaker  
Cardiff University, UK  
Tel: +44 29 20876999  
E-mail: [R.M.Whitaker@cs.cardiff.ac.uk](mailto:R.M.Whitaker@cs.cardiff.ac.uk)

## Owner-Centric Networking: A New Architecture for a Pollution-Free Internet

by Claude Castelluccia and Mohamed Ali Kaafar

*The PLANETE research group at INRIA Grenoble in France is developing an 'owner-centric networking' architecture. This novel concept will considerably reduce 'data pollution' and improve privacy on the Internet.*

In his essay (IEEE Security & Privacy, January/February 2009), Bruce Schneier famously said that data is the pollution of the Information Age. Content on the Internet (documents, emails, chats, images, videos etc) is often disseminated and replicated on different peers or servers. As a result, users lose the control and ownership of their content as soon as they release it.

The crux of the problem is that the Internet simply never forgets, and information that is posted lingers virtually forever. Furthermore, the design of the current Internet places no limit on data diffusion, nor any right to an individual to modify or remove what he/she wrote on a forum chat, or on a famous social network's walls.

This data pollution creates many privacy concerns, since this lost content can be used to collect information about users without their consent. For example, there have been several recent cases of employers using social networks (such as Facebook) to spy on their employees.

The Internet of the Future should solve these data pollution and privacy problems. However, according to Schneier, "Privacy isn't something that occurs naturally online, it must be deliberately architected". More specifically, we argue that the future Internet should give individuals control over their data. Users should be able to retrieve their previously posted content in order to withdraw or modify it. In other words, the Internet should enforce the 'right to

forget', which is a constitutional law in several countries.

Unfortunately, most if not all future Internet architecture proposals seem to have ignored this issue so far. For example, the content-centric networking (CCN) architecture, which proposes that the focus be shifted from transmitting data by geographic location to disseminating it via named content, actually increases data pollution. In CCN, content is not only hosted by servers but also diffuses from its point of creation to where the consumers are. As a result, individuals completely lose control over their content as it becomes distributed (lost) on the Internet without their consent or even knowledge.




Figure 1: A General Overview of the OCN architecture.

That said, we believe that content-centric networking is still a very attractive solution, if it then evolves towards an owner-centric architecture (OCN) that considers content ownership as bedrock. The OCN architecture that we are proposing comprises two main phases.

The first is content distribution and control. As in CCN, content diffuses to where the consumers are. However, in contrast to CCN, it is stored or cached in places controlled by owners and not by the network, such that it can easily be retrieved. These places should of course be defined according to the users' requests and needs, but be under the owner's control. All content belonging to an individual is under its owner's control; conceptually, it is as though a given individual's content is linked with a rope. Data are distributed and stored on the Internet, but at any time, users can pull on their ropes in order to retrieve them. Each time new content is added, a new element (a knot) is added to the rope. This rope can, for example, be implemented with a distributed hash table mechanism.

The second phase is content access. A user who wants to access content cannot download it (unless she owns it), but is only authorized to access it via a link, as occurs today when a user browses the Internet. As a result, instead of storing all the content locally, a user only stores the links to the content (except for that which she owns). As an illustration of this concept, an email becomes a collection of links on peers where the content of the email can be read. Similarly, a chat history is only composed of links to the actual conversation contents.

We believe that our OCN architecture would considerably reduce pollution on the Internet and improve privacy by giving users control over their data.

This project is still in its preliminary phase and we are aware that there are still many open issues to be solved. In particular, our scheme does not solve all privacy issues. For example, it does not prevent a service provider (such as search engines) from collecting information about its users for profiling or business purposes. These issues can be solved by integrating into this new

architecture concepts borrowed from anonymizing networks, such as onion routing and hidden services.

Furthermore, our scheme is only an architecture proposal, and as such does not prevent malicious users from violating it; for instance, by copying contents instead of just accessing them remotely. However, we believe that this last issue can be mitigated with the help of security protocols (such as SPKI certificates) and enforced by laws. Note that the situation is very similar to environmental pollution. Technology can help to reduce environmental pollution but cannot enforce it. For example, no technology can prevent a boat from emptying its fuel tank in the ocean and polluting the water. Only legislation and law enforcement can help in such a case.

In conclusion, we believe that the new paradigms developed in this project should be the subject of more attention and debate within the community. We advocate that future Internet architecture proposals should consider data pollution from the beginning. Privacy on the Internet is probably as important as security, and deserves equal consideration.

**Links:**

<http://www.schneier.com/essay-253.html>  
<http://mags.acm.org/queue/200901/?pg=3D8>

**Please contact:**

Claude Castelluccia  
 INRIA, France  
 Tel: +33 4 76 61 52 15  
 E-mail: ccastel@inrialpes.fr

Mohamed Ali Kaafar  
 INRIA, France  
 Tel: +33 4 76 61 55 95  
 E-mail: kaafar@inrialpes.fr

## Semantic Web for e-Commerce

by Bernd Gruber

*Imagine that your computer could understand the kinds of data it handles and how to interlink them in an intelligent way: that is what the Semantic Web is all about. ebSemantics is establishing this innovative technology in the Austrian tourism industry to enable Web users to efficiently find relevant information on the Web.*

The Federal Economic Chamber of Austria, the Austrian National Tourist Office, the e-business standardization

initiative AUSTRIAPRO and Smart Information Systems are working together to create the prerequisites for

taking advantage of e-commerce based on the Semantic Web. In particular, we foresee that the Semantic Web should

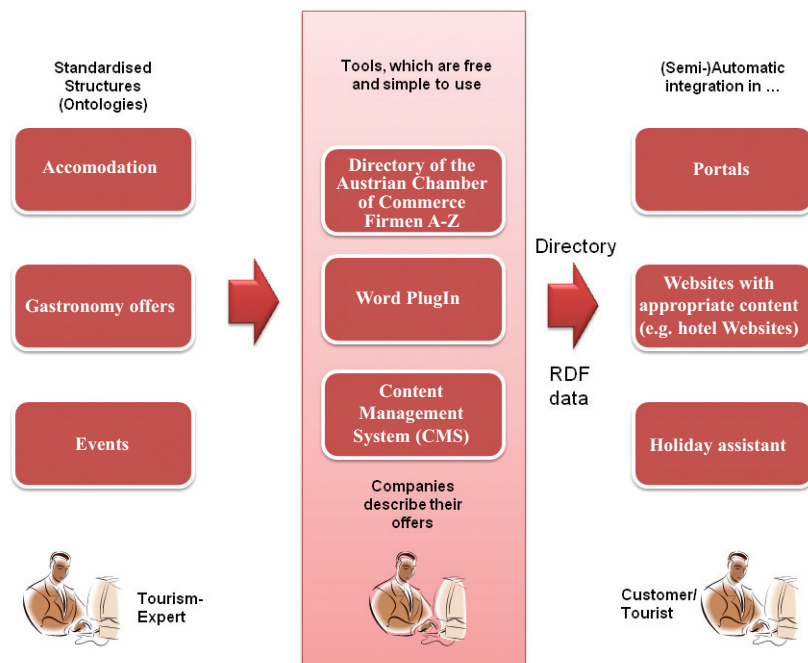


Figure 1:  
ebSemantics project sequence.

be easily accessible by small and medium-sized businesses with few resources.

To reach this goal, entry barriers must be minimized. It is therefore part of this project to develop an infrastructure for the use of tourism businesses, as well as creating a demonstration tool to illustrate through practical examples the advantages of this technology. In addition, the project will provide best practise guides based on the experiences made in the use cases.

Currently, while conventional search engines are able to search for keywords on Web sites, only human users can read and interpret product and service information on the Web. The innovative technology used by the Semantic Web enables intelligent applications and search engines to grasp the meaning of the information, process this data and display it so that it is comprehensible to humans.

This 'next-generation Web' is currently in a developmental stage. ebSemantics is taking action to develop semantic technology for the Austrian economy.

So what is necessary to make the Semantic Web work for Austria's e-commerce?

First, there must be standardized structures, so-called ontologies, which should be easy to implement in various applications. To achieve this aim, Smart

Information Systems and the University of Innsbruck are developing the collaborative wiki-based 'ontology development platform', myOntology.

Next, free and easy-to-use tools must be made available to everyone. An example of such a tool is the event platform OpenEvents, which publishes events on the Semantic Web. The Firmen A-Z (FAZ), a company register system that includes the basic data of more than 400,000 Austrian companies and is provided by the Austrian Federal Economic Chamber, is another way in which commercial offers and information can be advertised to the public. The FAZ is in fact one of the world's first Semantic Web-based directories that provides trusted authentication. In cooperation with Microsoft Austria, a Microsoft Word plug-in has been created to easily describe events and offers.

The final step is to interlink all the data, such as event information, catering offers and accommodation. As a result of this intelligent cross-linking between data, many diverse use cases occur. For example, a hotel Web site may contain information about relevant music events, or can link to a map containing restaurant locations. These events or restaurant offers can be customized so that, for example, only premium events are shown on a five-star hotel homepage or only cyclist-friendly restaurants are recommended along bicycle routes. In general, auto-

matic and intelligent integration, combination and connection of offers is possible.

The design underlying the Semantic Web is very simple. Based on OWL (Web Ontology Language) files provided by the research platform myOntology, it is possible to create RDF (Resource Description Framework) files, which contain information about a service. OWL is a family of knowledge-representing languages for ontologies. RDF was originally designed as a metadata model. It developed into a general method for conceptual description or modelling of information, which is implemented in Web resources. Both RDF and OWL are families of World Wide Web Consortium (W3C) specifications.

Future activities will include dissemination work such as attending fairs and congresses. Support and information will be provided to small and medium-sized businesses, leading to approval of and demand for the Semantic Web and improving the chances of technology uptake.

#### Links:

<http://www.ebsemantics.net>  
(only available in German)  
<http://www.ebsemantics.net/doc/>

#### Please contact:

Bernd Gruber  
Smart Information Systems, Austria  
Tel: +43 664 833 35 64  
E-mail: [b.gruber@smart-infosys.net](mailto:b.gruber@smart-infosys.net)

# Management of Future Mesh-Based Radio Access Networks

by Vangelis Angelakis, Vasilios Siris and Apostolos Traganitis

*Wireless mesh networks enable denser coverage and higher data rates than traditional wireless local area networks (LANs), while significantly lowering the deployment and operation costs compared to 3G mobile networks. These cost reductions combined with symmetric bandwidth for uplink/downlink communication will make wireless mesh networks a key architecture for future IP-based radio access networks, enabling innovative pervasive services based on ubiquitous broadband access.*

Wireless mesh networks combine three key features: the use of wireless transmission, the multi-hop nature of the end-to-end communication, and the availability of multiple paths. While each of these features raises a host of technical issues, at the same time they provide those advantages that make wireless mesh networks appealing. For instance, wireless transmission offers ubiquitous coverage for both stationary and mobile users; multi-path communication provides dependability in terms of increased network availability and reliability of wireless access to the Internet; and wireless multi-hop communication enables operation with reduced interference and low management costs. The overall mesh network management consists of the initial configuration, and subsequently adaptation, of procedures and mechanisms to achieve operation of the network according to some performance objectives. In this direction, we are investigating and developing new solutions for key network management functionalities, which include the following:

- power and rate control in metropolitan wireless mesh networks
- channel assignment in metropolitan wireless mesh networks with directional antennas and mesh networks with omni-directional antennas
- topology control in wireless mesh networks with omni-directional antennas.

Key features of our approach include the investigation of the cross-layer interaction of various mechanisms, the investigation of the timescales of management functions, and the combination of analysis and simulation, together with testbed implementation and evaluation.

The experimental testbeds include a fourteen-node indoor network with mesh nodes containing four radio interfaces, and a metropolitan mesh network. The latter covers an area of approxi-

mately 60 km<sup>2</sup> in the city of Heraklion in Crete, Greece and contains fourteen nodes, of which six are core mesh nodes, each containing up to four radio interfaces. The mesh network has two

case of low-quality links. For power control, we find that the minimum transmission power to achieve a target minimum performance does not change significantly during the course of a day.



*Figure 1: The experimental metropolitan wireless multi-radio mesh testbed in Heraklion, Crete. The mesh testbed contains IEEE 802.11a links with lengths ranging from 1.6 to 5 km. Nodes K1-6 are core mesh nodes, while nodes Mx are used for management and monitoring.*

fixed network gateway nodes: one at the Foundation for Research and Technology - Hellas (FORTH) and one at the University of Crete. Another important feature of the metropolitan testbed is an independent overlay network that is used for management and monitoring.

With respect to rate and power control in metropolitan wireless mesh networks with directional links, our results indicate that there are no performance gains when these mechanisms are performed on small timescales, of the order of packet transmission times. Indeed, adaptation of the transmission rate in very small timescales can actually reduce the achievable throughput in the

In the area of channel assignment in metropolitan wireless mesh networks, we have developed a practical procedure that accounts for both the interference between links belonging to the same mesh network (intra-network interference), and the interference due to sources external to the network (external interference). The former is captured through a Multi-Point Link Conflict Graph, which is appropriate for networks with directional links that include point-to-multipoint and multipoint-to-multipoint links. On the other hand, external interference is captured through measurement-based channel selection metrics, such as the signal-to-noise ratio or the packet delay.

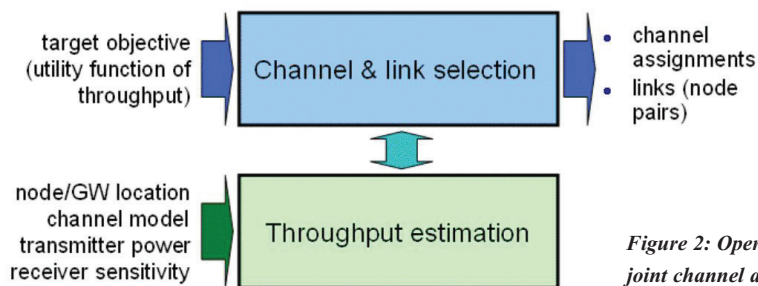


Figure 2: Operation and interaction of the modules for joint channel assignment and topology control.

Similar to the results concerning the timescale of rate and power control, we find that for metropolitan links using IEEE 802.11a technology, there are no significant performance gains to be had by performing channel assignment in short timescales, of the order of hours or even days. These results suggest that changes that affect the mesh network's performance do not occur on small timescales. On the other hand, unexpected events such as anomalies and intrusion attacks should be addressed through appropriate detection and mitigation mechanisms.

In metropolitan wireless mesh networks with directional links, the node connectivity or network topology is known a priori. In contrast, the network topology of wireless mesh networks with omnidirectional antennas is not. For such networks, we have developed a framework for joint channel assignment and

topology control. The corresponding channel assignment and topology control procedure consists of two modules: the throughput estimation module and the channel and link selection module (see Figure 2).

The innovative feature of the proposed framework and procedure is the support for different target objectives, which are expressed as functions of the achievable throughput between mesh nodes. The support for different target objectives is motivated by different operator-dependent requirements in terms of the aggregate throughput, the fairness of the link throughput and the link redundancy.

This research was conducted within the 7th Framework Programme project EU-MESH (Enhanced, Ubiquitous, and Dependable Broadband Access using MESH Networks), ICT-215320.

#### Links:

<http://www.ics.forth.gr/tnl>  
<http://www.eu-mesh.eu>

#### Please contact:

Vasilios Siris  
 FORTH-ICS, Greece  
 Tel: +30 2810 391726  
 E-mail: [vsiris@ics.forth.gr](mailto:vsiris@ics.forth.gr)

## Testing Mobile Data Applications on Smartphones with SymPA

by Almudena Díaz and Pedro Merino

*Due to its computing power, portable information, good connectivity, proximity to users, multimedia hardware such as cameras, and localization hardware such as GPS, the smartphone is positioned as the primary connection device for the future Internet.*

Our work focuses on the testing of data communication protocols for smartphones over cellular networks. The performance of communication protocols will play an important part in the evolution of mobile data services. New and more demanding services are expected for the future Internet, requiring new protocols for content distribution (multihoming protocols allowing the ample range of communication technologies

embedded in mobile devices to be exploited), new connectivity strategies (distribution of resources due to the high mobility of users) and new application requirements (multimedia content, point-to-point, real time etc).

The energy consumption of these new services is worth noting. Smartphones being mobile devices, battery life needs to be optimized in order to safeguard

their main functionality: to be in contact with the rest of the world, anytime and anywhere.

In that complex context, preparing for the deployment of new mobile services requires more than simulations. To evaluate and test factors such as usability, accessibility, availability and performance, we need to take into account real implementation and deployment



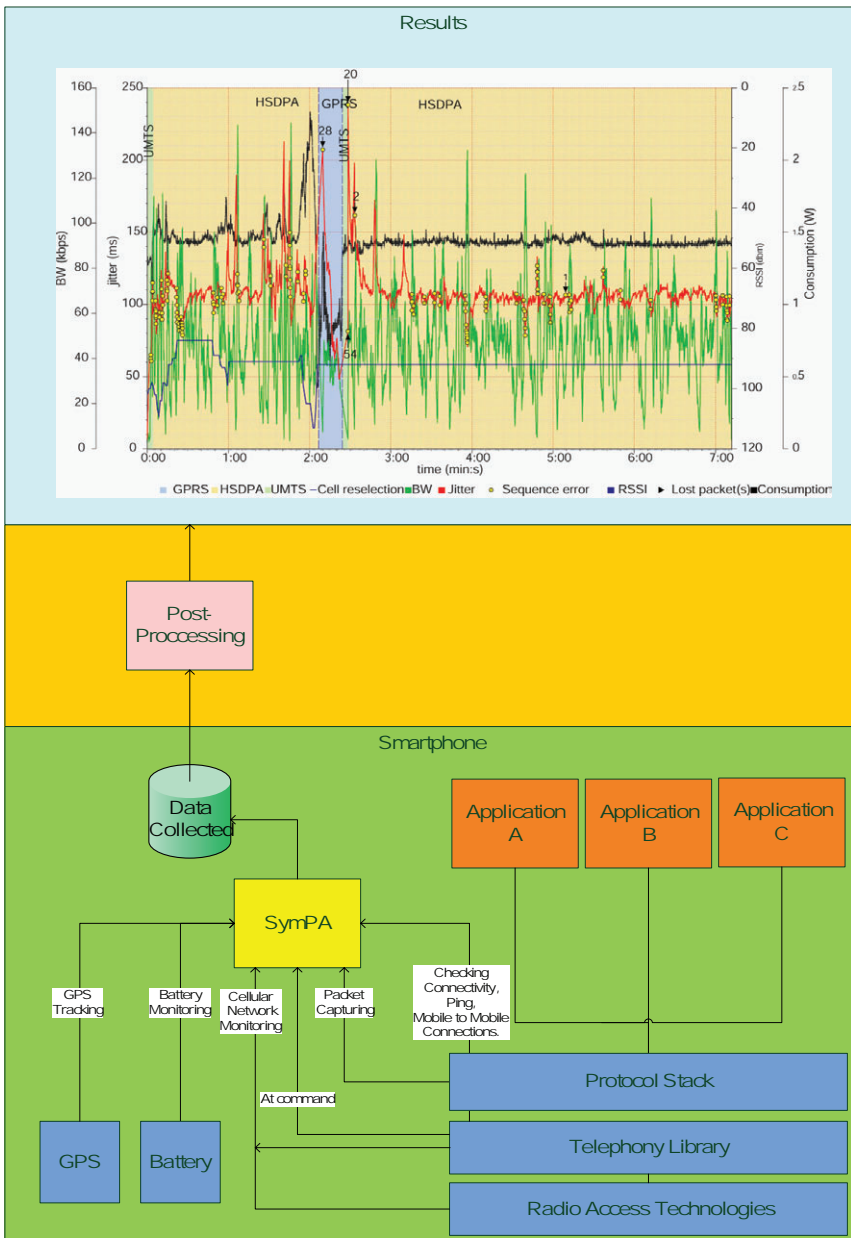


Figure 1: Performance of SymPA and the results obtained during post-processing of the collected data.

constraints, real users and real mobility patterns. Realistic field trials are thus required.

To cover all these issues, we have developed SymPA (Symbian Protocol Analyzer). SymPA is a software tool that runs on commercial Symbian OS smartphones, so third-party developers can use it to debug their applications in the real world. It can also be used to analyse the behaviour of third-party mobile applications and the performance of cellular networks from the point of view of mobile users.

Packet capturing is SymPA's main functionality. It enables the analysis of mobile protocols implemented by new

applications and of the performance of different data flows over cellular networks. Data traffic is correlated with information obtained by monitoring on the cellular network the radio access technology in use, quality-of-service profile parameters, received signal strength and cell information (network identifier, cell identifier, country code, location area code etc). That correlation helps us to identify issues such as changes in the radio mode (GSM, WCDMA, HSDPA, HSUPA etc), cell reselections or link outage, and their effect on IP data traffic.

Monitoring the battery consumption means the power profile of a given application and its relation to the radio

access technology in use can both be characterized. This functionality is crucial to optimizing applications for battery life.

All this information is accompanied by location data obtained from the GPS tracking function, which enables geographical positioning of the collected data. Finally, SymPA incorporates active functionalities that check the IP connectivity of the mobile device and characterize the round-trip times of data connections and direct mobile-to-mobile connections. An AT command interface is also provided.

Figure 1 illustrates the performance of SymPA and the results obtained during post-processing of data collected by it. The graphic represents three kinds of information. Traffic data such as jitter, delay, lost packets and sequence error have been obtained by processing packet capture with a network protocol analyser. The figure's background colour represents the technology in use, highlighting the burst of lost packets that takes place when there is a handover between HSDPA and GSM networks. It is also evident that power consumption increases when the received signal strength decreases before the handover to the GSM network, and that it decreases in GSM. This kind of correlation between data traffic, radio measurements and battery consumption enables us to characterize data communication performance over cellular networks.

This tool can also be used by operators who can take advantage of its functionalities to detect underperformance issues in their networks or to analyse the quality of service perceived by final users. The analysis of the data traffic received by the user's terminal allows the bandwidth perceived by customers to be measured, along with the real impact of radio issues on a user's experience.

**Link:**  
<http://www.lcc.uma.es/~pedro/mobile>

**Please contact:**  
 Almudena Díaz  
 University of Malaga/SpaRCIM, Spain  
 Tel: +34 952 132846  
 E-mail: [almudiaz@lcc.uma.es](mailto:almudiaz@lcc.uma.es)

# Fighting Lymph Cancer with Mathematics

by Annette Kik and Gunnar Klau

*What does a disease like lymph cancer have in common with city heating? This unusual association popped up when CWI researcher Gunnar Klau discussed his work with a fellow researcher. Klau now exploits maths originally developed for optimizing heating networks to analyse data from cancer patients. His method appears to be faster and better than previous ones, and it is hoped that his results will help doctors in understanding the disease and in the end contribute to the quest for new and better medicines.*

The types of cancer studied by Klau and his colleagues are called non-Hodgkin lymphomas - various malignant cancers originating in white blood cells and spread by the lymph node system. The disease is rather common, with several hundred patients in the Netherlands being diagnosed each year. In order to ensure a good prognosis and provide the best treatment, it is important to recognize the specific subtype of cancer that is involved. These cancer subtypes can be distinguished by a different kind of gene expression, ie the way genetic information is converted into functional products like proteins.


In molecular biology, the focus of research has recently shifted from decoding genetic sequences to analysing interactions that take place among genes and proteins. This is part of a new approach called systems biology. It is important to understand these interactions, since proteins do not function in isolation. They interact with each other and with other biomolecules to form molecular machines. These interactions are described by protein interaction networks. Huge amounts of experimental data on these networks are emerging from new high-throughput techniques like micro-arrays. Finding in a big network an active part (a subnetwork) that might contribute to the disease is like finding a needle in a haystack. When two thousand proteins are involved in a network, up to  $4 \times 10^{6595}$  candidate subnetworks must be investigated.

In spite of the apparent impossibility of this task, Gunnar Klau found a new method with which to tackle it. How does it work? "Compare it to city heating," Klau explains. "You can have big houses and small houses, far away or close by. Big houses that are located close by are most profitable for the energy company. Due to their short lengths, the pipelines are not too expensive. On the other hand, a small house

far away is not profitable and will probably not be connected to the city heating network at all. Traditional optimization methods only sought the big houses nearby and stopped calculating the moment they noticed a less favourable connection. However, it is possible that

a few minutes to enumerate possible interesting subnetworks!

Klau and his fellow researchers tested his system with known medical data, and found not only familiar subnetworks but even possible new interesting



*Lymphoma cancer cells. Picture: Dr. Lance Liotta Laboratory.*

just beyond the small house a group of big, profitable houses can be found. Our method can find these kinds of groups."

In Klau's metaphor, a group of big houses stands for an active protein subnetwork. The problem of finding these subnetworks is very difficult: in mathematical terms it is known as 'NP complete'. The lymphoma networks studied in this research count about 2000 nodes and 8000 connections. Klau used methods from discrete optimization to avoid looking at the exponential number of all subnetworks. First, he made a mathematical transformation to a known problem, the Prize Collecting Steiner Tree (PCST) problem. Then, together with statisticians from Würzburg, he developed a toolkit called 'Heinz'. Using data from medical experts, it took Heinz only

ones. The method is already better and faster than existing ones. "We still have to improve its accuracy and robustness," Klau says. "High-throughput biological data are extremely noisy, and the hidden subnetwork signals are quite weak."

In future, Klau wants to improve the model by integrating additional data sources to find biologically even more meaningful subnetworks. Klau: "This influences the mathematical model, because also the connections in the network will receive a score, based on co-expression of the two connected proteins. Then, the elegant transformation to the PCST-problem will not work anymore. We deal – mathematically – with a totally new problem, for which own theory has to be developed".

In the long term, it might be possible to analyse specific patient data faster with Klau's method. The pilot study has shown that it can distinguish between two types of lymphoma cancer. It might also recognize healthy patients from ill ones. "There is still a long way to go before this research can be applied in hospitals," Klau says. "However, I expect that it can help in studying network properties, identifying disease-related subnetworks, and network-based disease classifications. Ultimately, I hope the subnetworks computed with my mathematical methods

will help to create new biological and medical knowledge which might lead to better cures."

For this research, Klau and his fellow researchers won the Outstanding Paper Award at the prestigious ISMB 2008 (16th Annual International Conference Intelligent Systems for Molecular Biology) in Toronto, Canada. He cooperates closely with medical research partners like the Netherlands Kanker Instituut (NKI). "CWI finds it important to do this kind of interdisciplinary research," Klau says. "We started a full Life Sci-

ences Research cluster in 2009, to give this research a boost. I enjoy being part of this development."

**Links:**

- [http://homepages.cwi.nl/~klau/pubs/heinz\\_ISMB\\_2008.pdf](http://homepages.cwi.nl/~klau/pubs/heinz_ISMB_2008.pdf)
- <http://homepages.cwi.nl/~klau/>
- <http://www.cwi.nl/lifesciences>

**Please contact:**

Gunnar Klau  
 CWI, The Netherlands  
 Tel: +31 20 592 4117  
 E-mail: [Gunnar.Klau@cwi.nl](mailto:Gunnar.Klau@cwi.nl)

## Measuring Digital Library Usage Using Network Traffic Analysis

by Jiří Šmerda and Radka Findeisová

*A group at Masaryk University has developed a method that creates comparable statistical reports on the use of heterogeneous digital libraries. It achieves this by analysing network traffic to selected digital library repositories. Such statistical reports are crucial, particularly to aid institutions in evaluating and optimizing their digital library portfolios.*


Many research institutions subscribe to various providers of digital libraries, the subscription fees for which are often substantial. Institutions must therefore evaluate which digital libraries are used most frequently, which should be used frequently (but aren't), and which organizational unit uses each particular library the most. The results of the evaluation are used to optimize the portfolio of digital libraries to which the institution subscribes.

Digital library providers usually offer their own detailed statistical reports. These are clearly very useful for analysing the utilization of the selected library on its own. However, several difficulties exist. First, each provider offers its reports in a different format, making it hard to compare usage values for different providers. The second problem appears especially in larger institutions. In many cases it is unnecessary for the entire institution to sub-

scribe to a given library: only certain organizational units will require access. The summary reports for the whole institution are therefore unhelpful, because larger institutions want to break the usage figures down by organizational unit.

We have developed a method to deal with these problems. It uses data on network traffic collected by a hardware probe. The probe is attached to the point at which the institution is connected to the Internet. It collects the network traffic going to and from all computers located in the institution's local network. The collected data are filtered according to the digital libraries we want to monitor, and the results are aggregated, visualized and collated into reports. We measure the amount of data transferred from the digital library servers, the number of connections and the number of unique IP addresses that are connected to the digital library servers.

This research is taking place in the Institute of Computer Science at Masaryk University in the Czech Republic, in collaboration with the Faculty of Informatics. It has resulted in an application called MyLibScope, which



*Figure 1: Dynamic mindmaps in MyLibScope analytical desktop application - expanded nodes and edges with more information about usage within selected faculty and digital library.*

measures usage figures of over a hundred digital library repositories used by computers in the university network, which includes about 10,000 computers.

The application consists of three parts. The server part runs in the background, communicates with the network probe, collects and aggregates filtered data, and prepares them for visualization. The Web part takes the data and creates reports, which are available on the Masaryk University library Web site. The third part comprises an analytical desktop application that uses advanced methods of visualization to combine dynamic graphs, tables, forms and statistical graphs. This enables users to see the communication between faculties and digital library servers in a dynamic mind map: faculties and digital library servers are shown as nodes, while edges

between nodes represent their communication. The size of the node illustrates the amount of data transferred to and from the node. The thickness of the edge shows the amount of data transferred during the communication between two nodes. The values are relative to the time period a user selects for analysis. A user can interactively open details about each node and edge and see detailed usage graphs for faculties or digital libraries.

This work is just the beginning. We look forward to analysing usage data over longer periods, seeing trends and anomalies, and thoroughly evaluating the use of digital libraries at Masaryk University. In the future, we also plan to enrich the application with other data sources like the university information system or to include reports from digital library providers.

The application MyLibScope is the result of a broader project, which aims to enhance work with digital libraries both for end users and library administrators. The internal Masaryk University project 'Digital Libraries at Masaryk University' started at the beginning of 2007. MyLibScope was developed and deployed at the end of 2008 and will run in pilot phase from the beginning of 2009. The application uses hardware probes by INVEA-TECH a.s. and is built on software technology by Mycroft Mind a.s.

**Please contact:**

Jiří Šmerda

Institute of Computer Science, Masaryk University, Czech Republic

Tel: +420 549497676

E-mail: smerda@ics.muni.cz

## Multi-Faceted Rating of Product Reviews

by Stefano Baccianella, Andrea Esuli and Fabrizio Sebastiani

*Researchers from ISTI-CNR, Pisa, are working on an automatic rating system for online product reviews based on an analysis of their textual content..*

Online product reviews are now available on a variety of Web sites, and are being used by consumers with increasing frequency in order to make purchase decisions between competing products. For example, according to a study performed on TripAdvisor (one of the most popular online review sites for tourism-related activities) of users of the TripAdvisor online booking system, 97.7% are influenced by other travellers' reviews, and of those, 77.9% use the reviews as an aid in choosing the best place to stay.

It is obvious, therefore, that there is a growing market for software tools that can organize product reviews and make them easily accessible to prospective customers. Among the issues that the designers of these tools need to address are: (a) content aggregation, such as pulling together reviews from sources as disparate as newsgroups, blogs and community Web sites; (b) content validation, as in filtering out fake reviews authored by people with vested interests; and (c) content organization, as in automatically ranking competing products in terms of the satisfaction of consumers who have already purchased the product.

We address a problem related to issue (c), namely rating, which involves attributing a numerical score of satisfaction to consumer reviews based on their textual content. This problem arises from the fact that while some online product reviews consist of a textual evaluation of the product and a score expressed on some ordered scale of values, many other reviews only contain a textual part. Such reviews are difficult for an automated system to manage,

especially when a qualitative comparison of them is needed in order to determine whether product x is better than product y, or to identify the best product of the lot. Tools capable of interpreting a text-only product review and scoring it according to how positive it is, are thus of the utmost importance.

Our work looks at the problem of rating a review when the value to be attached must range on an ordinal (ie discrete)



*Example of 'product' rating.*

scale. This scale may be in the form either of an ordered set of numerical values (eg one to five stars), or of an ordered set of non-numerical labels (such as 'poor', 'good', 'very good', 'excellent'). We also focus on multi-faceted rating of product reviews, where the review of a product (eg a hotel) must be rated several times according to several orthogonal aspects (eg cleanliness, location etc).

We focus on generating the vectorial representations of the reviews that must be given as input to the learning device used to generate a review rater, rather than on the learning device itself (for which we use an off-the-shelf package). These representations cannot simply consist of the usual 'bag of words' used when classifying texts by topic, since classifying texts by opinion (which is the key content of reviews) requires a much subtler approach. Two expressions such as "A great hotel in a horrible

town!" and "A horrible hotel in a great town!" would receive identical 'bag of words' representations despite expressing opposite opinions.

We have focused on three aspects of the generation of meaningful representations of product reviews: (i) the extraction of complex features based on speech patterns; (ii) making the extracted features more robust through the use of a lexicon of opinion-laden words; and (iii) the selection of discriminating features through techniques explicitly devised for ordinal regression (an issue which until now has received practically no attention in the literature). In order to test the techniques we have developed, we crawled the Web to create a dataset of hotel reviews. The dataset is now available to the research community for experimentation. Several experiments that we have run on it confirm that a combination of these three techniques provides

the best performance on this particular type of data.

The system we have realized could work as a building block for other larger systems that implement more complex functionality. For instance, a Web site containing product reviews whose users only seldom rate their own reviews could use our system to learn from already rated reviews how to rate the others; another Web site containing only unrated product reviews could learn to rate its own reviews, from the rated reviews of some other site.

**Link:**

<http://nmis.isti.cnr.it/sebastiani/Publications/ECIR09c.pdf>

**Please contact:**

Fabrizio Sebastiani  
ISTI-CNR, Itlay  
Tel: +39 050 3152 892  
E-mail: [fabrizio.sebastiani@isti.cnr.it](mailto:fabrizio.sebastiani@isti.cnr.it)

## More on Mobile Telephones and Our Health

by Harry Rudin

*For well over two years, it has been hoped that the international investigation known as the Interphone Study would shed much-needed light on the question of whether mobile phone usage is a threat to our health. While the final results are not yet in, an excellent evaluation of the current status of the study has just been published. The indications are that if there are health dangers, they likely come from intensive, long-term use.*

The title of this comprehensive, well-written and easily digestible report is 'Comments on the Interphone Study'. It is available in German, French, Italian and English and was written by Gregor Dürrenberger, Jürg Fröhlich and Heinz-Gregor Wieser in January 2009. The authors are with, respectively, the Swiss Research Foundation on Mobile Communication, the Laboratory for Electromagnetic Fields and Microwave Electronics - ETH Zurich, and the Department of Neurology, University Hospital Zurich, all in Switzerland. The report was published under the auspices of the Research Foundation for Mobile Communication and the Krebsliga Schweiz, a Swiss organization for cancer research and education.

The topic is controversial, delicate and complicated, to the point that the many researchers involved in the Interphone Study have been unable to reach an



*Is mobile phone usage a threat to our health?*

overall conclusion. Many individual national studies have been published however, and the cited report describes these results with remarkable clarity. The hope remains that there will eventually be a single report which will

combine the many national studies, thus including more individual investigations and so leading to a more robust result. Statistically speaking, individually the national studies examine an insufficient number of cases to reach a solid conclusion. The reasons for this are also explained in the report.

The goal of the Interphone Study was to examine the potential link between mobile phone use and the risk of developing four different kinds of tumours in the head of the user. The approach used was epidemiological, using case-control studies. People were identified who had one of the four kinds of tumours and an investigation was made into their mobile phone use. The results were compared with those of demographically similar people who had not developed a tumour. The goal was to establish a connection between mobile phone use and the development of these tumours.

## Results

Since the individual national studies involved relatively small numbers of people, the certainty of the conclusions, statistically speaking, is questionable. As an example, in the case of two tumour types, there is (very weak) statistical evidence that mobile phone use decreases the risk of tumour development, a rather counter-intuitive conclusion. Another problem is the difficulty that study participants had in accurately estimating their mobile phone usage several years in the past.

Tumours take in the order of tens of years to develop, meaning there is great interest in 'long-term users'. The Interphone Study has its own definition of long-term mobile phone users: these are people with a 'regular' use pattern (at least half an hour per week) over a period of ten years or more. For these users there is an indication that longer exposure leads to greater risk based on the observation that tumours tend to develop on the side of the head where the user holds his or her phone. The difficulty is that the study overall has only a small number of users and an even smaller number of long-term users. Thus the statistical strength of the conclusion is still weaker in the case of long-term users.

It is worth noting, however, that today's younger generation spends hours per day with their mobile phones pressed to their heads, exposing themselves to much more radiation than the 'regular' users investigated in the study.

The brevity of this summary prevents us from exploring many interesting details, and the reader is encouraged to study the sixteen-page report itself (see link below).

## An Aside

Quite independently of the Interphone Study and other studies of mobile-phone use relating to cancer, there have been a number of studies relating to mobile-phone use and automobile accidents. Here the statistical certainty is much greater. The results show that using a mobile phone while driving (independently of whether or not it is a hands-free phone) increases the chances of having an accident by a factor of four: about the same increase as driving with a blood alcohol concentration at the legal limit! One such recent study, published in 2006, comes from Professor David Strayer and colleagues at the University of Utah (see link below).

## Conclusions

The mobile phone is a useful, convenient, even life-saving device, but it should be used responsibly and carefully:

- keeping the phone away from your body is a good idea; this can be achieved by using a headset
- putting the phone into stand-by or sleep mode reduces exposure to a negligible level
- avoiding prolonged conversations reduces exposure.
- using a mobile phone in a metallicity enclosed space such as a car, bus or

elevator (or wherever reception is weak) tends to increase exposure (since the phone generally increases its radiated power to compensate for poorer reception)

- chances are that children are more endangered by mobile-phone radiation than adults because their skulls are more permeable to electromagnetic fields. Hence if dangers exist, they might well be greater for children
- driving while carrying on a mobile-phone conversation is dangerous both for the user as well as for other people on the road.

## Links:

Interphone study:

[http://www.mobile-research.ethz.ch/var/Comment\\_Interphone\\_update01.pdf](http://www.mobile-research.ethz.ch/var/Comment_Interphone_update01.pdf)

University of Utah study:

<http://www.psych.utah.edu/AppliedCognitionLab/HFES2006.pdf>

## Please contact:

Harry Rudin

ERCIM News local editor for Switzerland

E-mail: [hrudin@sunrise.ch](mailto:hrudin@sunrise.ch)

# The Open Wall: A Software-Intensive Art Installation

by Letizia Jaccheri

*The Open Wall is a computer-based art installation developed by students, researchers and artists. The Open Wall enables its developers and observers to engage in dialogue that leads to increased understanding about creativity, cooperation, openness and authorship of software-intensive processes and content.*

The SArt project is conducted by the software engineering group in the Department of Computer Science at the Norwegian University of Science and Technology. We study the intersection of software and art in order to improve our knowledge of creativity, coopera-

tion, openness and authorship in software-intensive processes and content. The Open Wall project is one of SArt's projects.

In 2005 a group of architecture students built an experimental house with

an LED façade; the house was intended to last for a year. The students asked the Department of Computer and Information Science (CIS) for help and cooperation. Hardware design was the most important task when the installation was built for the first time. When

the experimental house was removed, the physical components were taken over by CIS. In 2007 a Masters student developed the Open Wall software into an open-source project with a BSD license. In January 2008 three groups of Masters students rebuilt the installation, reusing the existing hardware and software and developing the missing pieces of software and the content to be displayed on the Wall. In autumn 2008, the ITovation competition was organized and attracted twenty international contributors.

The Open Wall is a wall-mounted LED piece consisting of 96 circuit boards (16 x 6 boards) containing 2400 orange LEDs with a 5cm spacing. The wall is



*School children interact with the Open Wall.*

480cm long and 180cm high. Each board has 25 LEDs on its surface, which emit light with 99 possible intensities. Each board has its own micro-processor, and power and Ethernet connection. Connection to the main controller device is established through a set of switches or hubs.

While the installation was first erected on the façade of the experimental house in Trondheim, the current Open Wall is in a meeting room in the Department of CIS. It can be accessed by a Web-based interface: this enables users to upload and see the content of the wall with an Internet browser. The content produced by students, and by artists and engineers for the competition, is also on the Web.

Technologists and researchers see the cooperation with artists as a source of inspiration and an opportunity to reflect

on technology and find inspiration for innovations. Artists want to engage in projects like this to explore the possibilities inherent in technology and in interaction with technical people and researchers. Students choose this project as part of their curriculum because they enjoy cooperating with other students from different backgrounds.

Technology gets old quickly. Technologists experience this inevitability as a source of frustration but also of motivation to continue learning about new technology. There are a number of questions worth asking. How important is the type and novelty of technology in a cooperation project between artists and technologists? Is the Open




Wall a piece of art or a tool for artistic expression? Taking the installation from the public space into a meeting room in a University department has consequences in this respect. If the installation could be regarded as a piece of art when in was the public space, has it become a technological prototype or tool now that it has been moved into a private space in a university? What roles do the Web interface and source code play with respect to the artwork?

In traditional art the identity of the artist is important; in this project in this project, different questions become important. Who is responsible? How can we attract and facilitate multidisciplinary participation in the development of such projects? How can we facilitate good software evolution by publishing it as open source? And is the BSD

license most appropriate licence for this project?

### Conclusions

The novelty of our work lies in the nature of the process (an open cooperation between IT researchers, architects, artists and students), its objectives (artistic rather than business goals) and its open architecture, which enables everybody not only to upload content and become the artist but also to change the software governing the installation. Moreover, the process of bringing technology from the inside of a small computer or mobile phone to a big wall makes people reflect on and question IT in new ways. Software becomes visible and tangible.



The novelty of the single component is not important. Students and artists who choose to work with the Open Wall in both curriculum-driven and hobby-driven projects do not pose questions about technology age. They are motivated by the unusual way in which the Open Wall puts components of technology together.

### Links:

SArt: <http://prosjekt.idi.ntnu.no/sart/>

Open Wall software:

<http://sart.svn.sourceforge.net>

Open Wall project documentation:

<http://www.theopenwall.no>

### Please contact:

Letizia Jaccheri

NTNU, Norway

Tel: +47 91897028

E-mail: [letizia@idi.ntnu.no](mailto:letizia@idi.ntnu.no)

# Simulating Invisible Dangers in Emergencies: The SimRad-Projects

by Gerhard Chroust, Karin Rainer and Markus Roth

*The importance of preparedness against chemical, biological, radiological and nuclear (CBRN) dangers is growing immanently. SimRad.NBC generates the foundation for a user- and practice-oriented simulation and communication tool for first responders in such emergency scenarios.*

In case of a CBRN emergency, first responders (ie fire brigade, emergency medical services and police) must evaluate the situation within minutes and take appropriate action to minimize negative effects on persons and property.

An important problem concerning chemical, biological, radiological and nuclear dangers is that we do not have any inborn, natural sensors with which to recognize most of them, let alone natural, semi-autonomous reactions. It is therefore necessary to provide tools for identifying these dangers. This includes training in how to operate the equipment, interpret measured data, draw correct conclusions and initiate appropriate actions. We must also identify, design, validate and train appropriate behaviour ('best practices') to counteract these dangers. Some best practices might be counterintuitive and, if not well chosen, might negatively interfere with one another (again an area for validation). The choice of the appropriate strategy will often depend on a systemic evaluation of the whole situation, including necessary long-term considerations. This means that in order to cope with such dangers and threats, both first responders and specialists need to be well trained under realistic and practical conditions, including human factors, reactions and inter-organizational communication and cooperation.

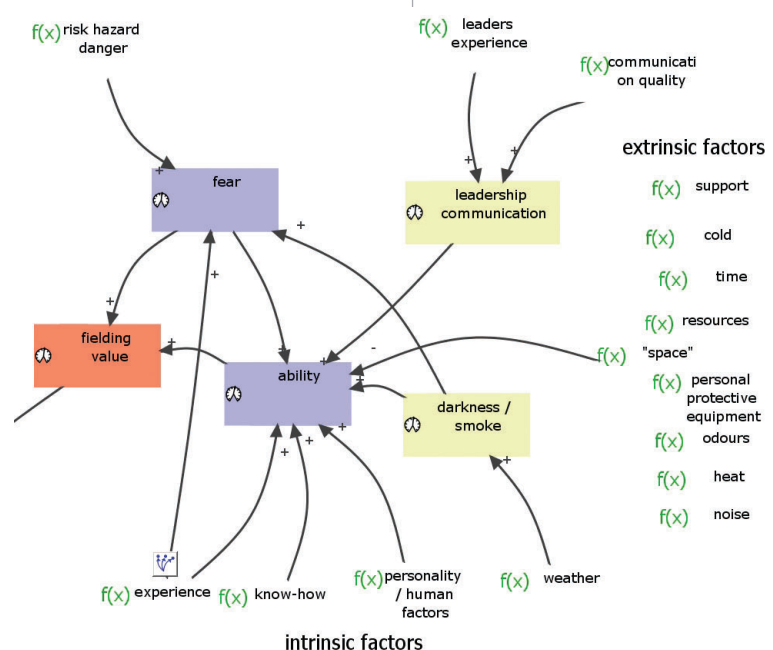
Until now, realistic training under 'safe conditions' (like flight simulators) was largely unfeasible; appropriate environments for the safe and realistic reproduction of such situations without creating real dangers were lacking. Modern technology allows both the actual situation and the consequences of various alternative actions to be simulated in a holistic way with real or simulated tools used in emergencies. Beyond the immediate training needs, simulations provide further advantages such as the easy repeatability of a training situation, recording of all details, testing of com-

patibility of best practices and the visualization of subtle influences. Near-realistic simulation also allows the inclusion of the best practices of experienced practitioners, who can comment according to the 'SIK'-principle ('Seeing Is Knowing'). Additionally, those simulations and tools can be used during real emergencies to predict and plan patterns and where possible to evaluate alternative actions.

Vital to the SimRad-project series is its focus on the needs and demands of the first responders themselves. Besides the scientific recherche of pre-existing international studies, a main source of relevant input for the technical side of research and development is the praxis-driven sight of the first responders. To pay tribute to this focus, the consortium of SimRad.NBC consists, along with two SME R&D partners and a university institute taking care of the systemic and holistic aspects, of the Schwechat Fire Department, the CBRN-Department of

the Red Cross Upper Austria and the Research Institute of the Red Cross.

Implicit in the core focus of SimRad.NBC on end-user requirements is its perspective on various positive and negative influences and connections involving technical equipment and the surrounding environment. These forms of influence are used to create a model of different crisis situations. The model - built on system dynamics principles - is designed to visualize relationships between influence factors with and without direct connection. This is the foundation for generating a powerful and above all practical simulation tool, which mirrors the complexities of a real CBRN-crisis. Factors in cause and effect systems such as motivation, risk perception, personal equipment and training experience, are identified and fed into loops or work as multipliers as well as silencers. The output of this model is a map of generic influences and their dependencies.



*Figure 1: SimRad.NBC simulates the complex interaction of various extrinsic and intrinsic influences on first responders.*



For this theory-centred part of the project work, the valuable views of the end users and their suggestions and ratings have been gathered in a series of interviews called SimRad.NEEDS. Experienced trainers as well as trainees involved in CBRN-response are asked for their ratings of emergency communication, influences on the first responders and current models of training and practice with respect to simulating 'invisible dangers'. New and astonishing findings reveal the right path to the analysis of similarities and differences between the different groups of first responders.

The current project SimRad.NBC began in 2007. On completion in November 2009, it will be immediately succeeded by SimRad.COMP. Based upon the foundations of its predecessor, SimRad.COMP will take further steps to generate feasible pre-prototypes for a simulation and communication tool-pack to tackle 'invisible CBRN-dangers'. Both projects are supported by the Austrian Federal Ministry for Transport, Innovation and Technology (bmvit) within the KIRAS Security Research Programme.

**Links:**

Project Web site: <http://www.simrad.at>  
KIRAS: <http://www.kiras.at>

**Please contact:**

Peter Ziehesberger  
Technical Project Leader  
Ziehesberger Elektronik, Austria  
E-mail: [peter@ziehesberger.at](mailto:peter@ziehesberger.at)

Stefan Schönhacker  
Research Institute of the Red Cross,  
Austria  
E-mail: [stefan.schoenhacker@w.rotekreuz.at](mailto:stefan.schoenhacker@w.rotekreuz.at)



## DL.org: A Coordination Action on Digital Library Interoperability, Best Practices and Modelling Foundations

by Donatella Castelli and Stephanie Parker

*What is playing an increasingly central role in the way we communicate, deal with information, learn, conduct commerce, work, deliver and use government services, carry out research, and practise and benefit from health care? The answer is: digital libraries.*

Digital Libraries (DLs) are making the vision for universally accessible collections of human knowledge a reality. For this to happen, considerable advances in digital library technologies and systems are needed. In particular, there is the need for Global Digital Library Infrastructures resulting from the federation of regional, national, international interoperable DL systems and digital repositories worldwide. However, making DL systems interoperable is a complex process that needs to take into consideration core requirements for digital library architecture, content, functionality, policy, quality and user perspectives.

DL.org, a recently funded EC project, which stems from the DELOS project, is paving the way for the future interoperability of DL systems thus making feasible the implementation of Global Digital Library Infrastructures. The DL.org Consortium includes: CNR – ISTI (Coordinator), University of Athens, HATII at the University of Glasgow, and Trust-IT.

DL.org aims to forge strong alliances with major stakeholders and the wider digital library community to harness global expertise and knowledge and

maximize opportunities for the enhanced development of the next generation of Digital Library systems. Key DL.org goals focus on capitalising key achievements, addressing the main open research & technical issues, capturing best practices, promoting shared standards, and producing guidelines for the development of global DL infrastructures.

DL.org coordinates a number of Working and Liaison Groups composed of representatives of international DL initiatives and internationally known experts and is organising eTraining courses designed to enable the next generation of DL professionals.

The DELOS Digital Library Reference Model will constitute the basic conceptual framework within which DL.org will conduct its activities addressing the DL interoperability problem from the Architecture, Content, Functionality, Policy, Quality and User perspectives.

Expected outcomes of DL.org include:

- A "Digital Library Technology and Methodology Cookbook" providing a portfolio of current best practices and patterns to facilitate cross-fertilization among existing systems and their

enhancement in terms of critical interoperability issues.

- A consolidated and enhanced version of the DELOS Digital Library Reference Model.
- Workshops, summer schools, eCourses and Dissemination activities to communicate the impact of DL.org achievements to relevant communities.

DL.org is also working with standardisation organisations like the Open Grid Forum and its European partner, OGF-Europe, to identify OGF standards that can support the interoperability goals of the project, as well as define further developments of them.

**Links:**

<http://www.dlorg.eu/>  
<http://www.delos.info/>  
[http://www.delos.info/files/pdf/ReferenceModel/DELOS\\_DLReferenceModel\\_096.pdf](http://www.delos.info/files/pdf/ReferenceModel/DELOS_DLReferenceModel_096.pdf)

**Please contact:**

Donatella Castelli  
ISTI-CNR, Italy  
Tel: +39 050315-2902  
E-mail: [Donatella.Castelli@isti.cnr.it](mailto:Donatella.Castelli@isti.cnr.it),  
[info@dlorg.eu](mailto:info@dlorg.eu)

# FLaReNet: Shaping the Future of a Multilingual Digital Europe

by Nicoletta Calzolari

*The launching event of FLaReNet - Fostering Language Resources Network - was held in Vienna, 12-13 February 2009, as 'The European Language Resources and Technologies Forum'. This event combined the FLaReNet themes with the i2010 objectives to address some of the technological, market and policy challenges to be faced in a multilingual digital Europe. The forum represented an important opportunity to identify the grounds for future directions in Language Resources and Technology at the level of the EC, national governmental agencies and industry.*



*Nicoletta Calzolari opens the FLaReNet Vienna Forum.*

The new EC eContentPlus Thematic Network (ECP-2007-LANG-617001) FLaReNet – Fostering Language Resources Network – aims at developing a common vision in the field of Language Resources (LRs) and Language Technologies (LTs) and fostering a European strategy for consolidating this area. The goal is to enhance competitiveness both in the EU and worldwide.

The growth of the field should be complemented by a common effort that identifies synergies and overcomes fragmentation. By creating consensus among major players in the field, the mission of FLaReNet is to identify priorities as well as short, medium and long-term strategic objectives and provide consensual recommendations in the form of a plan of action for the EC, national organizations and industry.

Through the exploitation of new collaborative modalities as well as workshops and meetings, the FLaReNet Network will sustain international cooperation and (re)create a broad-based Language Technology community.

Work in FLaReNet is inherently collaborative. A set of Working Groups cover various thematic areas and carry out their activities through workshops, meetings and via a col-

laborative Wiki platform. The FLaReNet Thematic Areas and Working Groups are:

- the chart for the area of LR and LTs in its different dimensions
- methods and models for LR building, reuse, interlinking, maintenance, sharing and distribution
- harmonization of formats and standards
- definition of evaluation and validation protocols and procedures
- methods for the automatic construction and processing of LR.

Anyone interested in these areas can subscribe to the FLaReNet website, join any of the working groups and participate in their activities. This will offer the advantage of playing a role in the definition of recommendations for future actions, thus shaping the future with respect to the new challenges.

The Forum was composed of a series of working sessions in which leading experts presented their visions on hot topics in the field of Language Resources and Technologies. In order to elicit new ideas and perspectives from the widest possible audience, a contest for the best contribution was held and the winner invited to give a presentation in the evaluation session. In all the sessions, discussants and participants actively contributed to the discussion and the current debate about priorities in the sector.

The final session was dedicated to a round table on international cooperation, mainly with non-European participants. Future policy and priorities were discussed in a global context, and the aim was to initiate a strategic discussion on the utility of promoting international cooperation among initiatives and communities around the world.

All position papers and slides can be found on the FLaReNet Web site. Proceedings of the event with reports on the outcomes of the various sessions, main trends and challenges will be available soon.

## Join FLaReNet!

FLaReNet is bringing together leading experts from many research institutions, academies, companies, consortia, associations, funding agencies, and European and international public and private bodies. In order to initiate the process of establishing a stakeholders' platform that is as pervasive as possible, the FLaReNet steering committee is inviting all interested players in the field to express their interest in becoming part of the network. To do this, fill in the form available on the project Web site.

In the meantime, stay tuned for the first issue of the FLaReNet newsletter. This is a means by which the network can reach the outside community and reinforce links between network subscribers. The first issue will focus on the FLaReNet launching event, featuring a report on the event as well as interviews with the invited speakers.

**Link:** <http://www.flarenet.eu/>

## Please contact:

Nicoletta Calzolari (FLaReNet Coordinator),

ICL-CNR, Italy

E-mail: [Flarenet\\_Coordination@ilc.cnr.it](mailto:Flarenet_Coordination@ilc.cnr.it)

## 1st Workshop on Formal Methods for Service-Oriented Architecture and Internet of the Future

The Working Group “Formal Methods for Service Oriented Architecture and Internet of the Future” run by the DEPLOY (Industrial deployment of system engineering methods providing high dependability and productivity) project held a one-day event on 16 February in Düsseldorf at the 7th international conference on integrated Formal Methods (iFM 2009). More than ten talks were presented on themes centering around Web services, SOA, fault tolerance, composition as well as using formal methods ranging from Event-B to rewriting logic. The workshop was very successful, with sessions being attended by 20 to 35 people. A new edition of the workshop is planned within the Formal Methods for Components and Objects (FMCO) symposium, which is part of the first Formal Methods Week (FMWeek) to be held in Eindhoven in November 2009.

This DEPLOY Working Group gathers expertise on Internet Technology and Formal Methods from FP7 Call 1 projects of Objective 1.2 “Services and Software Architectures, Infrastructure and Engineering”. A main objective of the Working Group is to contribute to the cooperation between these projects, in which several ERCIM members are participating (for example CWI, SARIT, SpaRCIM, and Fraunhofer Gesellschaft).

**More information:**

[http://www.formal-methods.de/ifm09/fmsoa\\_workshop.html](http://www.formal-methods.de/ifm09/fmsoa_workshop.html)

<http://www.deploy-project.eu>

<http://www.win.tue.nl/fmweek>

Call for Papers

## CSCLP 2009 Annual ERCIM Workshop on Constraint Solving and Constraint Logic Programming

Barcelona, Spain, 15-17 June 2009

The workshop will cover all aspects of constraint and logic programming, including various foundational issues, implementation techniques, new applications as well as teaching issues. Particular emphasis is on assessing the current state of the art and identifying future directions.

We would like to invite authors to submit papers covering all aspects of research on constraint and logic programming. Standard research papers, position papers and work-in-progress papers describing current projects are all welcome.

To submit a paper to the workshop, please use:  
<https://www.easychair.org/login.cgi?conf=csclp2009>

Papers, preferably in Springer LNCS format, should not exceed 15 pages. All submissions must be received by 8 May 2009.

The workshop is open to all and is not restricted to ERCIM members. There is no registration fee. Participants are kindly asked to register for the workshop before the registration deadline by sending a message to the workshop chair at [larrosa@lsi.upc.edu](mailto:larrosa@lsi.upc.edu) including your name, affiliation, e-mail and phone number.

**More information:**

<http://www.lsi.upc.edu/~larrosa/CSCLP2009.html>

Call for Participation

## The Service and Software Architectures, Infrastructures and Engineering (SSAIE) Summer School

Crete, Grece, 16-19 June 2009

The Transformation Services Laboratory (TSL) of the Computer Science Department of the University of Crete in cooperation with the Institute of Computer Science (ICS-FORTH) organizes a Summer School on “Services and Software Architectures, Infrastructures and Engineering (SSAIE)”. In addition to high quality training, the Summer School helps forge a new research and scientific community on Service Science Management and Engineering (SSME). The Summer School fosters the free exchange of ideas and helps the participants to network and start new cooperative research projects. The school will be held in Heraklion, Crete from 16-19 June 2009.

The Summer School is organized in close cooperation with the Training and Summer School Collaboration Working Group (CWG), that brings together the S-Cube Network of Excellence and projects of the IST SSAIE unit. It is a continuation of a series of workshops on Service Science and Engineering, the first of which was held in Crete, in May 2007. The topic of the first workshop was: “The Business Process in the Service Science”. The second workshop was organized by S-Cube and Engineering and held in Palermo, Italy from 2-6 June 2007. The topic was “Perspectives on Services”.

**More information:** <http://www.ssme2009.tsl.gr/>



## TrebleCLEF Summer School on Multilingual Information Access

Santa Croce in Fossabanda, Pisa, Italy, 15-19 June 2009

The aim of the Summer School is to give participants a grounding in the core topics that constitute the multidisciplinary area of Multilingual Information Access (MLIA). The School is intended for advanced undergraduate and post-graduate students, post-doctoral researchers plus academic and industrial researchers and system developers with backgrounds in Computer Science, Information Science, Language Technologies and related areas. The focus of the school will be on "How to build effective multilingual information retrieval systems and how to evaluate them".

### Programme

The programme of the school will cover the following areas:

- Multilingual Text Processing
- Cross-Language Information Retrieval
- Content and Text-based Image Retrieval, including multilingual approaches
- Cross-language Speech and Video Retrieval
- System Architectures and Multilinguality
- Information Extraction in a Multilingual Context
- Machine Translation for Multilingual Information processing
- Interactive Aspects of Cross-Language Information Retrieval
- Evaluation for Multilingual Systems and Components.

### Location and Dates

The Summer School will be held in the beautiful ex-convent Santa Croce in Fossabanda, Pisa. Santa Croce provides the perfect setting for study and discussions.

### Accommodation and Registration

A maximum of 40 registrations will be accepted. Tuition fees are set at 200 Euros up to 30 April and 350 Euros after this date. Tuition fees cover all courses and lectures, course material, lunch and coffee breaks during the School, the welcome reception on the evening of Sunday 14 June, and the Social Dinner on Monday 15 June. Accommodation will be on the School site at Santa Croce in Fossabanda.

### Financial Support for Students

A number of grants will be made available by TrebleCLEF and by the DELOS Association covering accommodation costs. Students wishing to receive a grant must submit a brief application (maximum 1 page) explaining why attendance at the school would be important for them. The application must be supported by a letter of reference from the student's advisor/supervisor or equivalent.

### More information:

<http://www.trebleclef.eu/summerschool.php>

13th european conference on digital libraries

# ECDL 2009

Sept. 27 - Oct. 02 2009 | Corfu, Greece

Conference Theme : Digital Societies

Sponsored by ERCIM

## 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, in 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).

### More information:

<http://www.ecdl2009.eu/>



Call for Papers

## Joint ERCIM Workshop on Software Evolution and International Workshop on Principles of Software Evolution

Amsterdam, The Netherlands, 24-25 August 2009

Anniversary edition – The Future of Software Evolution

- 20th Anniversary of ERCIM
- 10th Edition of the International Workshop on Principles of Software Evolution
- 5th Edition of the International ERCIM Workshop on Software Evolution

This workshop, co-located with the 7th joint ESEC/FSE Conference, is a merger of the annual ERCIM Workshop on Software Evolution and the International Workshop on Principles of Software Evolution. The rationale for a common event is to capitalize on the synergies to be found when theorists and practitioners meet.

### Theme and Topics

For this year's anniversary edition, the special theme of the workshop is "The Future of Software Evolution". We particularly welcome submissions that take a historical perspective on a particular facet of software evolution research, practice or education, argue for the current challenges, and speculate on what the future might bring. Other topics sought include, but are not limited to:

- Application areas: distributed, embedded, real-time, ultra large scale, and self-adaptive systems, Web services, mobile computing, information systems, etc.
- Paradigms: support and barriers to evolution in aspect-oriented, agile, component-based, and model-driven software development, service-oriented architectures, etc.
- Technical aspects: co-evolution and inconsistency management of various software artefacts, impact analysis and change propagation, dynamic reconfiguration and updating, architectures and notations for evolvability, etc.
- Formal aspects: Theories, principles, and models of software evolution, adaptation of evolutionary concepts and measures from other disciplines, etc.
- Managerial aspects: effort estimation and risk analysis for software evolution, processes explicitly supporting evolvability, etc.
- Empirical studies: quantitative or qualitative studies of the evolution context (organisations), of how developer teams, code and other artefacts evolve, etc.
- Practical aspects and industrial applicability of any of the above

Advertisement

## introducing ATLANTIS PREMIUM

Maximum visibility for your  
Conference and Workshop  
Proceedings

fast  
on-line  
professional  
reliable  
&  
freely available for everyone



ATLANTIS PRESS  
AMSTERDAM-PARIS

For more details, please check:  
<http://www.atlantis-press.com/proceedings/>

### Submissions

Four types of submissions are possible: short position and tool papers (up to 5 pages) and long research and industrial papers (up to 10 pages). Position papers may present wild and speculative ideas, and will be judged on the potential to generate interesting discussions at the workshop. Position papers addressing the workshop's theme are especially welcome. Tool papers should describe tools supporting software evolution. The workshop will include one or more sessions for tool demos. Research papers should describe innovative research, while industrial papers may describe the application or adaptation of known solutions to industrial case studies, and reflect on lessons learnt.

### Important Dates:

Abstracts due: 23 May 2009  
Papers due: 28 May 2009  
Notification: 19 June 2009  
Camera-ready copy due: 26 June 2009

### More information:

<http://ssel.vub.ac.be/iwpse-evol/>

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

Eindhoven, The Netherlands, 2-3 November 2009

The aim of the 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 are intended to bring together scientists and practitioners who 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.

These workshops also strive to promote research and development for the improvement of formal methods and tools for industrial applications.

### Topics

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 (e.g., 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.

### Invited Speakers:

- Dino Distefano, Queen Mary, Univ. London, UK
- Diego Latella, ISTI-CNR, Italy
- Thierry Lecomte, ClearSy, France
- Ken McMillan, Cadence, USA

FMICS 2009 is part of the first Formal Methods Week (FMweek), which will bring together a choice of events in the area. For the latest information on FMweek, see <http://www.win.tue.nl/fmweek>.

### More information:

<http://users.dsic.upv.es/workshops/fmics2009/>

### Call for Papers

## ADT2009 - First International Conference on Algorithmic Decision Theory

Venice, Italy, 21-23 October 2009

ADT2009 is an Interdisciplinary forum on Uncertainty and Robustness in Decision Making, Preferences in Reasoning and Decision, Decision Theoretic Artificial Intelligence and Learning and Knowledge Extraction

The COST Action IC0602 and the EURO Working Group on Preferences are proud to announce the 1st International Conference on Algorithmic Decision Theory. A new unique event aiming to put together researchers and practitioners coming from different fields such as Decision Theory, Discrete Mathematics, Theoretical Computer Science and Artificial Intelligence in order to improve decision support in the presence of massive data bases, combinatorial structures, partial and/or uncertain information and distributed, possibly inter-operating decision makers. Such problems arise in several real-world decision making problems such as humanitarian logistics, epidemiology, risk assessment and management, e-government, electronic commerce, and the implementation of recommender systems.

Contributions to the conference are sought in areas including:

- Unconventional Uncertainty Distributions
- Robustness in Decision Making
- Adversarial Risk Analysis
- Extreme Events Risk Management
- Preference Modelling and Aggregation
- Computational Social Choice
- Compact Representation of Preferences
- Preferential Queries in Data Bases
- Universal Languages for Preferences
- Multiple Criteria Decision Analysis
- Planning and Sequential Decision Making
- Algorithmic Efficiency
- Multi-Agent Systems
- Argumentation and Explanation in Decision Support
- Preferences and Policies Learning
- Knowledge Extraction
- Algorithmic Tuning
- Recommender Systems
- Decision Support Systems.

### Important Dates

- Paper submission: May 3, 2009
- Notification: June 14, 2009
- Final version: July 5, 2009

### More information:

<http://events.math.unipd.it/adt2009/>

## New EUROGRAPHICS Chair again from an ERCIM Institute

Roberto Scopigno, leader of the Visual Computing Lab of ISTI-CNR, Pisa, Italy, has been elected Chair of the Euro-



*Roberto Scopigno.*

graphics Association for the period 2009-2010. Roberto is a member of the Executive Committee of Eurographics since 2001 and has served as Vice-Chair for the last four years.

Eurographics (European Association for Computer Graphics) is an international organization dedicated to the needs of professionals, scholars and students in computer graphics and related visual disciplines. Roberto will follow the path traced by the previous Chair, David


Duce, Oxford Brookes Univ., aimed at raising the international profile and strengthening the association. In fact, Eurographics has played an increasingly important role in the scientific scenario in recent years, mainly through the organization of a stream of well-attended annual events, and the Eurographics publications are now considered landmarks by our computer graphics community. The aim of the new Chair will be to contribute to the further improvement of the scientific value of the association's initiatives and to ensure that Eurographics remains well tuned to the members' needs and expectations.

<http://www.eg.org/>

## New Book on Software-Intensive Systems and New Computing Paradigms

"Software-Intensive Systems and New Computing Paradigms – Challenges and Visions", edited by Martin Wirsing, Jean-Pierre Banâtre, Matthias Hözl and Axel Rauschmayer, and published by Springer in the "Lecture Notes in Computer Science" series vol. 5380, presents the results of the INTERLINK Working Group on software-intensive systems and novel computing paradigms.

INTERLINK is a 'Coordinated Action', managed by ERCIM and funded by the Future and Emerging Technologies (FET) Unit of the European Commission, to identify emergent trends, their impact on the information society in the next 10-15 years, and the challenges they present to computing, software engineering, cognition and intelligence. The objective of the working group was to imagine the landscape in which the next generations of software-intensive systems



will operate. To this end three workshops were organized on this topic to which over 30 leading researchers from Europe, Asia, Australia, USA and Canada were invited to present and discuss future R&D directions, challenges, and visions in this emerging area.

This volume starts with an overview of the current state of the art and the research challenges in engineering software-intensive systems. The remainder of the book consists of invited papers of the working group participants and is structured in three major parts: ensemble engineering, theory and formal methods, and novel computing paradigms. These papers cover a broad spectrum of relevant topics ranging from methods, languages and tools for ensemble engineering, socio-technical and cyber physical systems, ensembles in urban environments, formal methods and mathematical foundations for ensembles, orchestration languages to disruptive paradigms such as molecular and chemical computing.

<http://interlink.ics.forth.gr/>

<http://www.springerlink.com/content/978-3-540-89436-0>

## Master of Science in Computational Biology at Nice Sophia Antipolis University

While biological data exhibits a formidable diversity, the past two decades have seen the advent of massive data produced either by high throughput experiments or by measurement devices of increasing accuracy at very different scales ranging from nano to macro. Handling these massive and complex data within a virtuous cycle linking modelling and measurements is one of the major challenges in Computational Biology.

To take up this challenge and foster new discoveries in and new applications in medicine, the ambition of Master of Science in Computational Biology is to provide a state-of-the-art training geared towards computational biology in computer science, applied mathematics and statistics to address biological and medical problems. So this new master program from Nice Sophia Antipolis University is specially tailored for students interested in a truly interdisciplinary learning experience.

The scientific goal of this program is to focus on the human being from different perspectives (understanding and modelling functional aspects or interpreting biomedical signals from various devices) and at different scales (from macro-level to micro-level).

The aim of this program is to provide excellent academic or industrial career opportunities by offering high level coverage of modelling and computing principles that will enable the challenges to be met and make tomorrow's technological choices in biological, medical computing domains. To achieve this, classes will be given by several outstanding professors and researchers from the research institutes present in the campus. The teaching language is English.

<http://www.computationalbiology.eu>



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



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



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



Irish Universities Association  
c/o School of Computing, Dublin City University  
Glasnevin, Dublin 9, Ireland  
<http://ercim.computing.dcu.ie/>



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



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



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



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



Centrum Wiskunde & Informatica

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



Universitas Varsoviensis  
Universitas Wroclaviensis

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/>



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



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/>



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



Spanish Research Consortium for Informatics and Mathematics,  
D3301, Facultad de Informática, Universidad Politécnica de Madrid,  
Campus de Montegancedo s/n,  
28660 Boadilla del Monte, Madrid, Spain,  
<http://www.sparcim.es/>



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

FRNS  
rue d'Egmont 5  
B-1000 Brussels, Belgium  
<http://www.frns.be/>



Swedish Institute of Computer Science  
Box 1263,  
SE-164 29 Kista, Sweden  
<http://www.sics.se/>



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



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



Fraunhofer Gesellschaft

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



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



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



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](mailto:office@ercim.org)**

I wish to subscribe to the

printed edition  online edition (email required)

Name: .....

Organisation/Company: .....

Address: .....

Postal Code: .....

City: .....

Country: .....

E-mail: .....

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