ERCIM NEWS

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

Number 68, January 2007

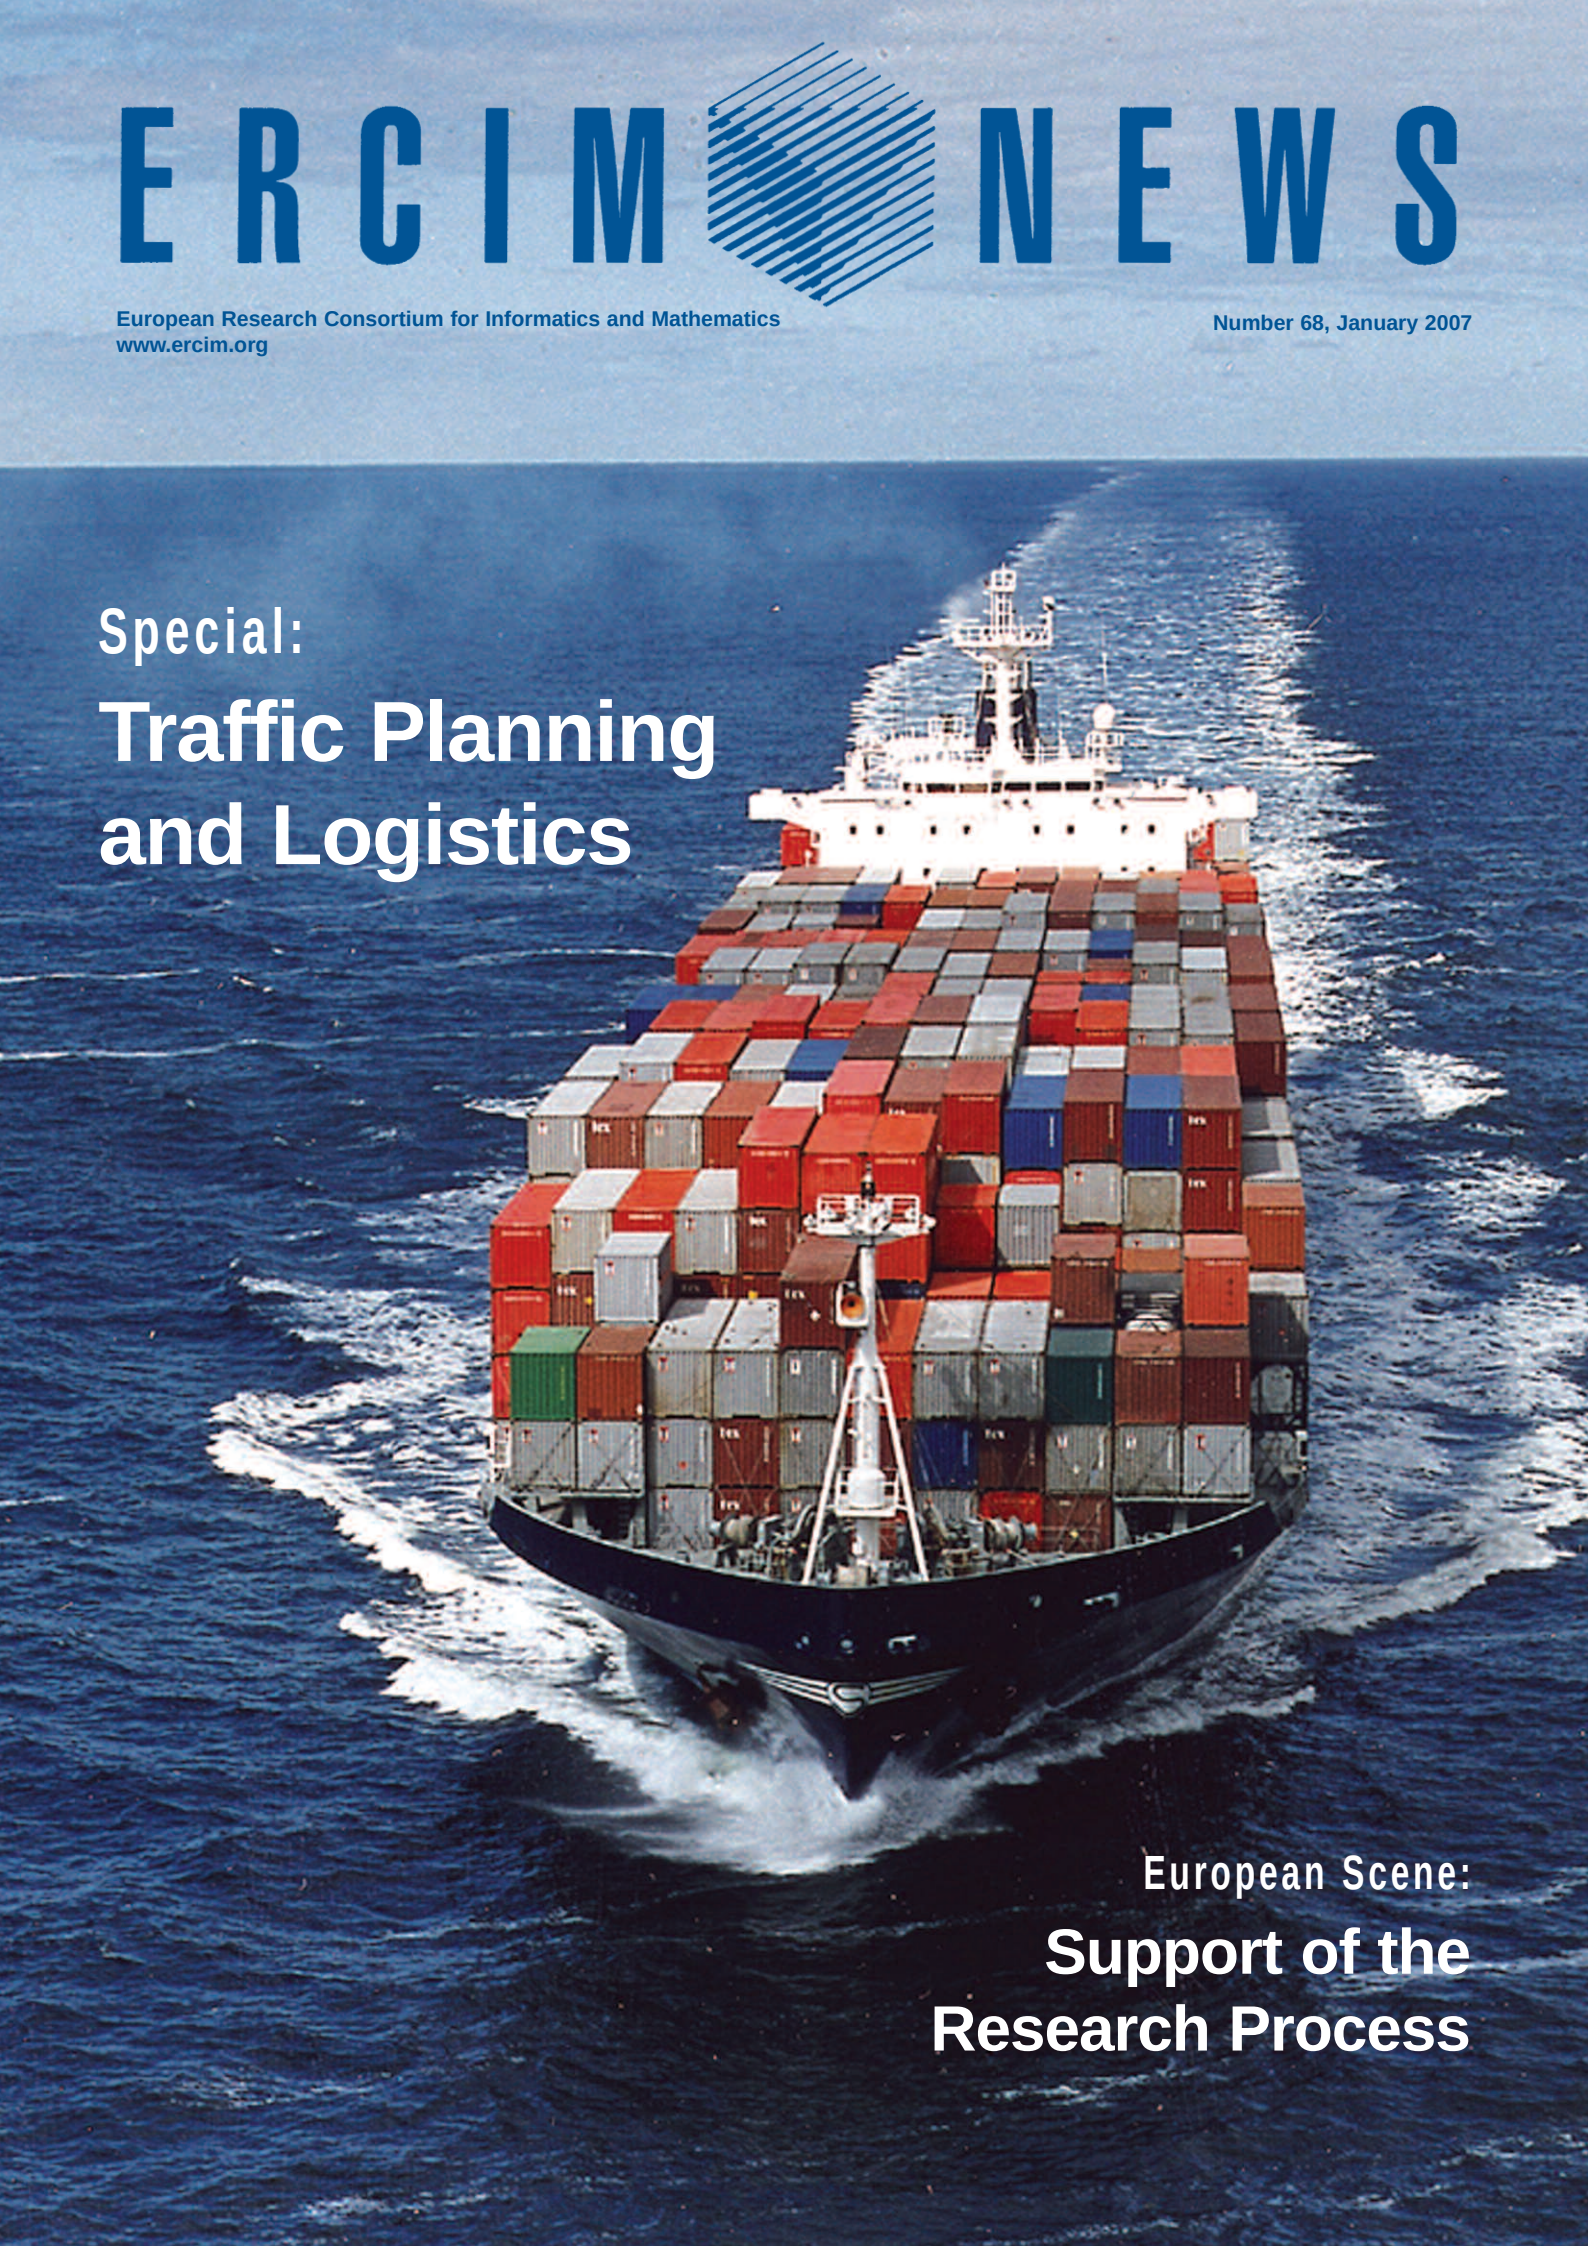

Special:

Traffic Planning and Logistics

European Scene:

Support of the Research Process

JOINT ERCIM ACTIONS

4 **InterLink: International Cooperation Activities in Future and Emerging ICTs**
by Constantine Stephanidis and Jessica Michel

5 **ERCIM at IST 2006 - En Route to FP7**

6 **The 3rd Grid@Asia and GFK 2006 International Joint Workshop**
by Bruno Le Dantec

6 **New ERCIM Working Group on 'Smart Environments and Systems for Ambient Intelligence'**
by Anthony Savidis and Norbert Streitz

7 **ERCIM Working Group 'User Interfaces for All': A Success Story Draws to a Close**
by Constantine Stephanidis

8 **9th ERCIM Workshop 'User Interfaces for All'**
by Constantine Stephanidis and Michael Pieper

9 **Cor Baayen Award 2007**

10 **MUSCLE CIS Coin Competition 2006**
by Allan Hanbury and Michael Nölle

11 **Workshop on Multidisciplinary Simulation in Biomedical Applications**
by Marc Thiriet

11 **New Projects managed by ERCIM**

NEWS FROM W3C

12 **W3C Celebrates Ten Years with Style**

13 **W3C Workshop on Web of Services for Enterprise Computing**

13 **Third W3C Workshop on Internationalising the Speech Synthesis Markup Language (SSML)**

13 **Markup Validator, Link Checker Updated**

13 **W3C10 Asia: Tenth Anniversary of W3C in Asia**

13 **Latest W3C Recommendations**

THE EUROPEAN SCENE 'SUPPORT OF THE RESEARCH PROCESS'

14 **The European Strategy Forum for Research Infrastructures**
by John Wood

16 **Europe: The Place to Be!**
by Peter Nijkamp

17 **CERIF - the Common European Research Information Format**
by Keith G Jeffery

18 **Managing Research and Knowledge Based Organisations**
by Günter Koch

SPECIAL THEME 'TRAFFIC PLANNING AND LOGISTICS'

20 **Introduction to the Special Theme**
by Jo van Nunen

22 **Mathematics for Railway Timetabling**
by Leo Kroon

23 **Railyard Shunting: A Challenge for Combinatorial Optimisation**
by Per Kreuger and Martin Aronsson

25 **The Art of Stacking**
by Martin Aronsson and Per Kreuger

27 **Freeway Applications Based on the Macroscopic Model of Traffic Flow**
by Tamás Luspay, István Varga and Balázs Kulcsár

28 **Performance Evaluation of Heterogeneous Motorized Traffic at Urban Single-Lane Junctions and Roundabouts**
by Puspita Deo and Heather J. Ruskin

30 **Fluid-Dynamic Approach to Traffic Flow Problems**
by Gabriella Bretti, Roberto Natalini and Benedetto Piccoli

31 **Dynamic Routing of Rail Vehicles**
by Anders Holst and Markus Bohlin

32 **Statistical Models for Urban Traveller Information and Traffic Management Systems**
by Poulicos Prastacos and Yiannis Kamarianakis

Cover photo: Schenker AG.

33 **Mobile Environmental Sensor Systems Across a Grid Environment - the MESSAGE Project**
by John Polak

34 **Statistical Physics Algorithms for Traffic Reconstruction**
by Arnaud de La Fortelle, Jean-Marc Lasgouttes and Cyril Furtlehner

35 **From Traffic Prediction to Collaborative Navigation Solutions**
by Arnaud de La Fortelle, Angel Talamona and Mikaël Kais

37 **Logistics Planning Using the POEM Language System**
by Jianyang Zhou, Jianyi Zhou, Jiande Zhou, Yann Courtet, Weiping Wu, Hua Ni

39 **Using GIS for Optimisation in Transportation Planning**
by Günter Kiechle

40 **GIS Technology for Maritime Traffic Systems**
by Cyril Ray, Thomas Devogele, Valérie Noyon, Mathieu Petit, Sebastien Fournier and Christophe Claramunt

42 **PeerMart: Decentralized Auctions for Bandwidth Trading on Demand**
by David Hausheer and Burkhard Stiller

44 **Net-WMS - A New Generation of Warehouse Management Systems Networked Services**
by Francois Fages and Abder Aggoun

45 **The IBM Secure Trade Lane Solution**
by Francois Dolivo

47 **Clipped RFID Tags Protect Consumer Privacy**
by Günter Karjoth and Paul Moskowitz

48 **TraSer – Identity-Based Tracking and Web Services for SMEs**
by Zsolt Kemény and Marcell Szathmári

50 **Distributed Engine for Advanced Logistics**
by Tamás Máhr and Alfons Salden

51 **Supporting City Planning with Traffic Monitoring in Image Sequences**
by Matthieu Molinier

52 **Traffic Planning and Logistics: A View from the Edge**
by Antonio Lucas

R&D AND TECHNOLOGY TRANSFER

54 **Software Travels in the Fast Lane: Good News or Bad?**
by Fabrizio Fabbrini, Mario Fusani and Giuseppe Lami

55 **Coping with System Complexity: Identifying** Dichotomic Architectural Alternatives
by Gerhard Chroust

56 **Power on Tap**
by Lars Rasmusson

58 **New Hungarian National Cancer Registry**
by László Kovács and Máté Pataki

59 **Efficiency and Quality of Service of Internet Service Providers**
by Oliver Heckmann

EVENTS

60 **CONCUR 2006 - The 17th International Conference on Concurrency Theory**
by Christel Baier and Holger Hermanns

61 **25 Years of the IFSR: A Milestone in the Systems Sciences**
by Gerhard Chroust

62 **PATAT 2006 Conference**
by Hana Rudová

63 **ESFORS – NESSI Workshop Addresses New Research Priorities in Trust, Security and Dependability in Software Systems and Services**
by Aljosa Pasic and Jim Clarke

64 **Announcements**

66 **Editorial Information**

66 **In Brief**

InterLink: International Cooperation Activities in Future and Emerging ICTs

by Constantine Stephanidis and Jessica Michel

In order to attract and foster trans-disciplinary research excellence, research programmes need to be defined around new grand challenges or key technological issues that are significant to economic growth or are derived from major societal drivers. Such programmes should be exploring visionary research themes that demand breakthroughs in basic research and engineering, and will be investigating radically new uses for technology.

Following the very successful Coordination Action 'Beyond the Horizon', the purpose of the newly launched EU Coordination Action 'InterLink' is to promote collaboration and partnerships in a number of strategic research areas in future and emerging Information and Communication Technology (ICT). InterLink aims at advancing Europe's knowledge in these areas, at promoting European solutions and knowledge worldwide, and at influencing the way in which research in these areas will evolve internationally. It will therefore carry out activities and provide means and support mechanisms to encourage cooperation between European and non-European research communities in strategic basic research areas related to Information Society Technology (IST).

Three thematic areas have been carefully selected. These are based on the need to address both the evolution of the information society in the next ten to fifteen years, and the challenges this will impose on research in computing, software engineering, cognition and intelligence. These thematic areas, which are also closely linked to ERCIM's areas of strength, are:

1) Software intensive systems and new computing paradigms: the continuing decrease in the cost of microprocessors and storage is leading to the development of increasingly distributed and decentralized systems. Society's dependence on software-intensive systems is increasing rapidly, both in the economic sector and in daily life. Applications will be assembled as dynamic federations of autonomous and evolving components. In this context, new computing paradigms and techniques are required to build software-intensive systems.

2) Ambient computing and communication environments: the evolution of the information society is characterized by the development of personalized individual and collective services. These exploit infrastructures situated in smart environments and are based on a range of ubiquitous and pervasive communication networks that provide ambient computing at multiple levels. The underlying vision of pervasive and ambient computing assumes very large numbers of 'invisibly' small computing devices embedded in the environment. Such devices would interact with multiple users in a wide range of dynamically changing situations, and be supported by an infrastructure of intelligent sensors (and actuators) embedded in the built environment. The realisation of this vision requires advances in various areas of ICT, and necessitates highly multidisciplinary research.

3) Intelligent and cognitive systems: cognitive systems should be able to interpret data arising from real-world events and processes, acquire situated knowledge of their environment, act, make or suggest decisions, and communicate with people on human terms. The design of cognitive artificial systems requires coordinated and integrated research efforts that span a wide range of disciplines and require collaborative resources that can be achieved only through long-term, multidisciplinary research activities.

The main aims of InterLink in the three selected thematic areas are:

- to identify and address world-scale, basic research problems of an emerging global information and knowledge society, where significant added value is expected to be gained from world-wide cooperation
- to establish communication and cooperation mechanisms within and beyond Europe in order to support the formation and functioning of a related scientific community
- to identify complementarities in the selected areas between EU and other involved countries that could give rise to exchanges of knowledge and technology
- to define joint basic research agendas, outlining research priorities and defining roadmaps and joint Research and Technological Development (RTD) initiatives.

InterLink is based on three working groups, one for each thematic area, that will work in a coordinated fashion. Working groups will consist of international teams of researchers with internationally recognized expertise in their fields of research. To support the Working Groups in conducting their work, advanced communication tools will be provided, aiming at a global networking of research communities. Several workshops will be organised during the project to foster RTD dialogue and promote interaction within each Working Group. Ultimately, InterLink will lead to the elaboration of a technology roadmap on an international scale in each thematic area.

The opening workshop for InterLink will take place from 10-12 May in the South of France (most likely in Nice, Antibes or Cannes). Detailed information will be available on the InterLink web site.

Link:
http://interlink.ics.forth.gr/

Please contact:
Constantine Stephanidis (scientific coordinator)
ICS-FORTH, Greece
E-mail: cs@ics.forth.gr

Jessica Michel (administrative coordinator)
ERCIM office
E-mail: jessica.michel@ercim.org

ERCIM at IST 2006 – En Route to FP7

The 'Information Society Technology 2006' conference and exhibition, organised in Helsinki in November 2006, was the opportunity for the European Commission to present its 7th framework programme (FP7) to the European IST research community.

The European Commission announced that it would invest over €9 billion in research on Information and Communications Technologies (ICT) within FP7 that will run until 2013. By allocating the largest single budget item in FP7 to this priority, the European Commission is acknowledging the strategic importance of ICT for Europe's growth and competitiveness.

Over 4,500 members of the IST research community were present in Helsinki to discuss the new framework programme and the strategic priorities for fundamental and applied ICT research throughout FP7.

ERCIM European Projects at IST 2006

Several ERCIM projects were presented during IST 2006, in particular CoreGRID and ACGT. These two leading IST initiatives had a dedicated booth in the main exhibition hall.

CoreGRID (the European Research Network on Foundations, Software Infrastructures and Applications for large scale distributed, GRID and Peer-to-Peer Technologies), considered as a model for Networks of Excellence (NoE), promoted its activities and highlighted a strategic shift in Grid technology development, with Grid computing research effort progressively moving towards the industry and industrial applications.

The ACGT project (Advancing Clinico Genomic Trials on Cancer) was hosted on the 'HealthGrid" booth. The project promoted the use of Grid technologies to support biomedical applications. A dedicated networking session was also organised to meet with other IST actors in e-health, with a special focus on grid applications to fight cancer.

Photo: European Commission.

General view of the IST 2006 exhibition.

Photo: European Commission

Viviane Reding, Commissioner for Information Society and Media during the opening keynote speech at IST 2006.

ENGAGE (Encouraging and Stimulating future RTD cooperation in the field of Information Society Technologies between Europe and Asia-Pacific region) was presented on the EU-Asia Cooperation booth. This initiative supports the development of deeper IST research collaboration with Asia and Southeast Asia.

Other ERCIM initiatives had organised specific and well attended networking sessions, in particular:

- MUSCLE: The Network dedicated to multimedia data mining and machine learning specifically presented its e-teams research integration effort (virtual labs) and its new showcasing activity destined to present MUSCLE technology to the industry.
- DELOS: The NoE on Digital Libraries presented its activities and plans concerning architecture development of digital libraries of the future; search in multiple languages on data bases; and collaboration on digital archives, which will be initiated with CVCE (Centre Virtuel de la Connaissance sur l'Europe, a public undertaking on Europe's heritage).

The IST 2006 conference has reinforced ERCIM's image as a key player in the IST research community, not only in Europe, but also worldwide.

ERCIM is now looking forward to new partnerships and projects in FP7, and seems to be particularly well positioned across a wide array of IST priorities.

Links:
IST 2006: http://europa.eu.int/information_society/istevent/2006/index_en.htm

ERCIM projects: http://www.ercim.org/activity/projects/

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

Third Grid@Asia and GFK 2006 International Joint Workshop

by Bruno Le Dantec

Some 150 representatives from European and Asian research institutes, industry and ministries of research participated in the Grid@Asia workshop in Seoul from 11-13 December 2006. The workshop was the third workshop organised by ERCIM to promote co-operation with Grid actors from South-East Asian countries in the European Union's 6th and 7th Framework Programmes (FP). This workshop was held in conjunction with the winter workshop of the Grid Forum Korea (GFK).

The workshop was opened by Jysoo Lee from the Korea Institute of Science and Technology Information (KISTI), followed by Joo-Han Kim from the Korean Ministry of Science and Technology who gave an overview of the cooperation status and driving plan for science and technology between Korea and the European Union. Jorge Gasos from

Photo: Ignacio Soler Jubert.

Won-Sik Kim, Deputy Minister, Ubiquituous Society Strategy Office, Ministry of Information and Communication, Republic of Korea, addresses the workshop .

the European Commission completed the opening session with a presentation of the the European Commission's Grid Technologies Unit. The workshop sessions of the first day addressed the role of Grid technologies in the 7th EU FP and some Grid research projects funded under the 6th FP were presented.

The second day was dedicated to specific Grid-related technologies such as Grid infrastructure, Grid systems and middleware with presentations of European and Korean projects. Satoshi Sekiguchi from the National Institute of Advance Industrial Science and Technology, Japan provided a keynote speech about Asian Grid projects. Speeches of Won-Sik Kim, Deputy Minister, Korean Ministry of Information and Communicaition, Byung-tae Yang, president of KISTI, and the presentation 'Past, Present and Future of Grid in Korea' given by Jysoo Lee, KISTI, Korea closed the second workshop day.

Grid applications were discussed the third day. Some 15 presentations addressed e-Science applications (ie, bio informatics, meteorology, etc) and business Grid applications (ie, onlime game services, ASP services, etc.). Presentations are available for download from the Grid Forum Korea and the Grid@Asia web sites

Link:
http://www.gridatasia.net/

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

New ERCIM Working Group on 'Smart Environments and Systems for Ambient Intelligence'

by Anthony Savidis and Norbert Streitz

The recently formed Working Group SESAMI focuses on research towards distributing, embedding, coordinating and interactively delivering computing intelligence, while putting people and social contexts at the centre of design considerations. This will result in smart surrounding environments and ubiquitous intelligence.

Ambient intelligence represents a vision of the near future in which 'intelligent' or 'smart' environments and systems react in an attentive, adaptive, and active (sometimes even proactive) way to the presence and activities of humans and objects in order to provide intelligent/smart services to the inhabitants of these environments.

Ambient intelligence technology integrates sensing capabilities, processing power, reasoning mechanisms, networking facilities, applications and services, digital content, and actuating capabilities distributed in the surrounding environment. While a wide variety of technology is involved, the goal of ambient intelligence is to hide its presence from users by providing implicit, unobtrusive interaction paradigms. People and their social situations, ranging from individuals to groups, be they work groups, families or friends and their corresponding environments (office buildings, homes, public spaces etc), are at the centre of the design.

The scope of the ERCIM Working Group SESAMI includes the facilitation of collaborations in this area, on the grounds of on-going, potentially cross-domain, basic and applied research and development. In this context, SESAMI will pursue novel insights on designing, implementing, managing

and maintaining smart computational environments on any scale. This will enable to enhance and go beyond traditional computational support of human activities for any given situation, context or role.

Background

Ambient intelligence and smart environments represent an emerging field of research and development that is rapidly gaining attention from an increasing number of researchers and practitioners worldwide. There are several reasons for this.

Ambient intelligence brings a special perspective to research associated with technical fields such as ubiquitous computing, pervasive and proactive computing, ambient computing, embedded computing and smart objects. In addition, the notion of ambient intelligence is becoming a de facto key dimension of the emerging Information Society, since many next-generation industrial digital products and services are shifting towards a smart computing environment.

Focus of Thematic Areas

Work related to ambient intelligence is concerned with modelling and supporting people and their activities, sensing and locating people and objects, managing and interpreting contexts, providing awareness and notification, and enabling people to interact naturally and intuitively with their environments via multiple modalities and coupled devices. Additionally, a key issue regarding smart environments is the delivery of extensible infrastructures with varying intelligent computational resources exploiting underlying information sources, corresponding software platforms and architectures able to provide and handle intelligent services. The orchestration and integration of computational elements that comprise distributed applications serving multiple users, both mobile and stationary, is another important aspect of ambient intelligence. Relevant applications are expected to provide value-added services via smart artefacts that are smoothly integrated with the surrounding environment and can also be accessed remotely in distributed environments.

Potential thematic topics of the SESAMI Working Group include, but are not limited to, the following:

- modelling and support of human activity
- context models, management and interpretation
- sensing paradigms and technology
- awareness and notification
- tangible, multimodal and implicit interaction
- universal access
- mobile and stationary artefacts
- distribution and orchestration
- service infrastructures
- multimedia and adaptation
- coordination and collaboration
- social aspects, reflecting the 'human in the loop'
- security and privacy
- applications, eg office/work, home, public spaces, health care, education and leisure/entertainment.

Current Activities

Two special thematic sessions are to be held as part of the 4th International Conference for Universal Access in Human Computer Interaction (UAHCI), jointly with the 12th International Conference on Human-Computer Interaction, on 22–27 July 2007 in Beijing, China (http://www.hcii2007.org/). These are SESAMI–I: Interaction Design and Evaluation, chaired by Norbert Streitz, and SESAMI–II: Tools, Architectures and Infrastructures, chaired by Anthony Savidis. An open call for papers has been made for both sessions, resulting in a total of fourteen accepted papers, with seven per session. All papers will be published as part of the Springer LNCS proceedings of the UAHCI 2007 conference.

A Special Theme issue will be published by Elsevier in their international, refereed, interdisciplinary journal 'Interacting with Computers' (IwC). The issue is entitled 'Interaction in Ambient Intelligence: Methods, Tools and Applications' and is guest-edited by Anthony Savidis, Norbert Streitz, Carsten Magerkurth and Constantine Stephanidis. The selection process for the papers has been initiated by the guest editors.

Link:
http://www.ics.forth.gr/sesami

Please contact:
Anthony Savidis, SESAMI Chair, ICS-FORTH, Greece
Tel. +30-2810-391749
E- mail: as@ics.forth.gr

Norbert Streitz, SESAMI Co-Chair
Fraunhofer IPSI, Germany
Tel: +49 6151 869-919
E-mail: streitz@ipsi.fraunhofer.de

ERCIM Working Group 'User Interfaces for All': A Success Story Draws to a Close

by Constantine Stephanidis

After twelve years of successful operation, the ERCIM Working Group 'User Interfaces for All' is drawing to a close. This Working Group has contributed significantly to the creation of a scientific community dedicated to Universal Access and Design for All in Europe, as well as to the establishment of this new research field.

Since its establishment in 1995, the ERCIM Working Group 'User Interfaces for All' (http://www.ui4all.gr/) has systematically promoted the Design for All principles in Human-Computer Interaction (HCI). The concept of 'User Interfaces for All' involves developing user interfaces for interactive applications and e-services that will provide universal access to interactive applications and services. The activities of the Working Group have included the organisation of a series of very successful workshops, which have helped consolidate recent work and stimulate discussion in the field. These workshops have brought together researchers from ERCIM institutions and other organisations worldwide, who share the common aspiration of making the information society equally accessible to all citizens.

The work of the ERCIM WG UI4ALL led in 2001 to the establishment of the International Conference on Universal Access in Human-Computer Interaction (UAHCI), which is part of the HCI International Conference series (http://www.hci-international.org/). UAHCI has established an international forum attracting participants from a broad range of disciplines and fields of expertise. Its purpose is to support the dissemination and exchange of results of theoretical, methodological and empirical research that addresses issues related to the attainment of universal access in the development of interactive software.

The Working Group has also contributed to the establishment of the international, interdisciplinary refereed journal 'Universal Access in the Information Society' (UAIS; http://www.springeronline.com/journal/10209/about), published by Springer. UAIS provides an archival publication channel for the discussion and advancement of theoretical and practical aspects of universal access in the information society, and stimulates cross-fertilization between the various contributing disciplines.

Many members of the ERCIM WG UI4ALL have also contributed to the edited book 'User Interfaces for All – concepts, methods and tools'. This is first book to have been dedicated to issues of universal design and universal access in HCI, and was published by Lawrence Erlbaum Associates in 2001 (http://www.ics.forth.gr/hci/publications/book.html, ISBN 0-8058-2967-9). It is a collection of thirty chapters written by leading international authorities who are affiliated with academic, research and industrial organisations and non-market institutions. It provides a comprehensive overview of the state of the art in the field, and includes contributions from a variety of theoretical and applied disciplines and research themes.

Furthermore, the work of the ERCIM WG UI4ALL has contributed to the establishment of the International Scientific Forum 'Towards an Information Society for All' (ISF-IS4ALL; 1997-2000), an international ad hoc group of experts who were the first to recognize the need for a global approach to an information society that is accessible, usable and acceptable by all citizens (http://www.ui4all.gr/isf_is4all/). Two White Papers have been published in the International Journal of Human-Computer Interaction and were also submitted to the European Commission, reporting on an evolving international R&D agenda in the field of HCI. Since then, the vision of an Information Society for All and the necessity for universal access to information society technology have acquired widespread acceptance and importance. This is the case not only at a scientific and technological level, but also at a European policy level, as demonstrated by the eEurope 'Information Society for All' initiative of the European Commission. The activities initiated by the ISF-IS4ALL continue in the framework of the Thematic Network (Working Group) 'Information Society for All' (IST-1999-14101 - IS4ALL, http://www.ui4all.gr/isf_is4all/), which has consolidated knowledge on universal access in the context of information society technology into a comprehensive validated code of design practice. This led to the publication in 2005 by Springer of the edited book 'Universal Access in Health Telematics - A Design Code of Practice' (http://www.springeronline.com/3-540-26167-2, ISBN: 3-540-26167-2). An online training course in Design for All (http://is4all-tc.ics.forth.gr) has also been developed.

As a result of the twelve years of activities of the UI4ALL Working Group, the field of 'Universal Access and Design for All' has made significant progress towards consolidating theoretical approaches, methods, tools and technologies, as well as exploring new application domains. The Working Group was the recipient of the ERCIM Working Group Award in 2000, and can be considered an example of a dedicated, prolific and successful research and development community that has contributed greatly to the establishment in Europe of a new scientific field.

As of the end of 2006, following the successful 9th Workshop, the time has come to consider the mission of this Working Group fully accomplished. At the same time, the emergence of ambient intelligence is raising new and fascinating research challenges, which are rapidly gaining attention from an increasing number of researchers and practitioners in Europe and worldwide. The notions of universal access and User Interfaces for All are central to this vision, since ambient intelligence aims at providing implicit, unobtrusive interaction by putting people, their social situations, and the corresponding environments at the centre of design considerations. These issues will be addressed by the new ERCIM Working Group 'Smart Environments and Systems for AMbient Intelligence' (SESAMI, http://www.ics.forth.gr/sesami/), which was established in 2006 (see article in this issue). The research community founded by the ERCIM Working Group UI4ALL will therefore continue its collaborative efforts in the context of SESAMI.

Link:
http://www.ui4all.gr/

Please contact:
Constantine Stephanidis
ICS-FORTH, Greece
ERCIM UI4All Working Group coordinator
Tel: +30 2810 391741
E-mail: cs@ics.forth.gr
http://www.ics.forth.gr/stephanidis/index.shtml

9th ERCIM Workshop 'User Interfaces for All'

by Constantine Stephanidis and Michael Pieper

In the tradition of its predecessors, the 9th ERCIM Workshop 'User Interfaces for All' aimed to consolidate recent work, and to stimulate further discussion on the state of the art in this area and its increasing range of applications in the Information Society.

The last of the workshops on 'User Interfaces for All' was held in Königswinter (Bonn), Germany, 27-28 September 2006. It was locally organised by Michael Pieper from

Fraunhofer-FIT, and built upon the results of the eight previous workshops, the first of which was held in Crete in 1995. Since 2001, the UI4ALL Working Group Workshop has alternated with the 'Universal Access in Human-Computer Interaction' (UAHCI) conference, which is held every two years in the context of the HCI International and affiliated conferences. In July 2007, Beijing will host the fourth UAHCI conference.

The emphasis of the 2006 event was on 'Universal Access in Ambient Intelligence Environments'. As a result of the increasing demand for ubiquitous and continuous access to information and services, it is anticipated that information society technology will evolve towards a new paradigm referred to as ambient intelligence. Such an environment will be characterized by invisible (ie embedded) computational power in everyday appliances and other common physical objects, and populated by intelligent mobile and wearable devices. Ambient intelligence will have a profound effect on the type, content and functionality of emerging products and services, and on the way people will interact with them. This will inevitably lead to multiple new requirements for the development of information society technologies.

In this process, the notion of universal access, which aims towards the accessibility and usability of Information Society Technologies by anyone, anywhere and at anytime – is critically important. The complex networks and interactivity involved in an ambient intelligence environment will present a number of challenges for universal access. The 9th ERCIM Workshop 'User Interfaces for All' therefore focused on these challenges, envisaging new and diverse scenarios of use of ambient intelligence technology, and identifying critical issues that will have to be addressed during the development life-cycle of interactive applications and services.

The Workshop offered two keynote speeches: 'From Human–Computer Interaction to Human–Environment Interaction: Ambient Intelligence and the Disappearing Computer' by Norbert Streitz (Fraunhofer IPSI, Germany); and 'Human Computer Confluence' by Alois Ferscha (Institut für Pervasive Computing, Johannes Kepler Universität Linz, Austria). Twenty-seven long papers and nine posters on topics related to ambient intelligence and universal access were presented during the Workshop, which attracted participation from all over the world.

The official Workshop proceedings will be published by Springer as part of the Lecture Notes in Computer Science series (LNCS 4397) and will be embedded in the LNCS digital library.

Link:
http://www.ui4all.gr/workshop2006/index.html

Please contact:
Constantine Stephanidis
ICS-FORTH, Greece
ERCIM UI4All Working Group coordinator
Tel: +30 2810 391741
E-mail: cs@ics.forth.gr
http://www.ics.forth.gr/stephanidis/index.shtml

2007 Cor Baayen Award

ERCIM calls for nominations for the 2007 Cor Baayen Award. Since 1995, ERCIM awards each year a most promising young researcher in computer science and applied mathematics with the Cor Baayen Award to honour Cor Baayen, the first president of ERCIM.

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

Conditions

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

Nominations should be made by a staff member of the university or research institute where the nominee is undertaking research. Self nominations are not accepted. A nomination must be submitted online (see 'Link' below). Deadline for nominations is 15 April 2007.

The 2006 Cor Baayen Award winner Oliver Heckmann receives the award from ERCIM President Keith Jeffery at a ceremony during the ERCIM fall meetings in November 2006. Read about Oliver Heckmann's work in the article 'Efficiency and Quality of Service of Internet Service Providers' in this issue (page 59).

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

Please contact:
László Monostori
SZTAKI, Hungary
ERCIM Cor Baayen Award coordinator
E-mail: laszlo.momostori@ercim.org

MUSCLE CIS Coin Competition 2006

by Allan Hanbury and Michael Nölle

A competition to find the best automated coin classification algorithm, with a prize sponsored by the MUSCLE Network of Excellence, was organised in 2006. The results were presented at a workshop held in September 2006 at the Heinrich Hertz Institute in Berlin in conjunction with the 28th Annual Symposium of the German Association for Pattern Recognition (DAGM). The prize of €1500 for the best performing coin classification program went to Marco Reisert, Olaf Ronneberger and Hans Burkhardt of the Chair of Pattern Recognition and Image Processing, University of Freiburg, Germany.

The training and test images for the competition are from the Coin Images Seibersdorf (CIS) database. This database was created as a result of the changeover from twelve European currencies to the Euro, when large volumes of mixed coin collections collected by charitable organisations had to be returned properly sorted to the national banks of the originating countries. During an automatic sorting process carried out at the Austrian Research Centers GmbH - ARC, the CIS database with manually verified classification of coin types was created.

Photo: Fraunhofer IPK Berlin.

Marco Reisert (centre) receives the prize from competition organisers Michael Nölle (left) and Allan Hanbury (right).

The training data provided to the participants consists of 30000 coins, corresponding to 60000 images (as the front and back of each coin is imaged). These are divided into 692 coin classes with 2270 different coin face classes. There are more coin face classes than coin classes as some coins changed their appearance over time, when new coin series were issued or designs to mark special occasions were introduced.

The task of each participant in the competition was to write a coin classification program and submit it to the organisers. Each program was then run on an unreleased competition dataset consisting of 10000 coins. Coins belonging to 362 of the 692 coin classes in the training data were present in the competition dataset, as well as 242 coins not belonging to any of these classes. The latter were to be classified as unknown. To simulate the demand by the national banks that only 1 in 10000 coins be classified incorrectly, an incorrect classification was penalised more heavily than a classification into the class of unknown coins.

Of the seven initial registrations for the competition, three programs were submitted. Two of these ran correctly on the competition data. The prize went to Marco Reisert, Olaf Ronneberger and Hans Burkhardt, University of Freiburg, Germany. Their program classified 9724 of the 10000 coins correctly and no coins incorrectly (all others were classified as unknown). The algorithm used is based on the analysis of the direction of gradient vectors.

More details on the competition, the algorithms submitted and the results can be found on the coin competition webpage. It is also possible to download the training set. The competition dataset, along with coin classifications, will be available here soon.

A second MUSCLE CIS Coin Competition is planned to take place in 2007. The details can be found on the coin competition webpage.

MUSCLE (Multimedia Understanding through Semantics, Computation and Learning) is a Network of Excellence managed by ERCIM.

Links:

MUSCLE NoE page: http://www.muscle-noe.org

MUSCLE coin competition page:
http://muscle.prip.tuwien.ac.at/coin_prize_2006.php

Chair of Pattern Recognition and Image Processing, University of Freiburg:
http://lmb.informatik.uni-freiburg.de

Seibersdorf Coin Recognition Machine:
http://www.smart-systems.at/products/products_image_processing_coin_en.html

Please contact:

Allan Hanbury
Technical University Vienna, Austria
Tel: +43 1 58801 18359
E-mail: hanbury@prip.tuwien.ac.at

Michael Nölle
Austrian Research Centers GmbH - ARC / AARIT, Austria
E-mail: michael.noelle@arcs.ac.at

Workshop on Multidisciplinary Simulation in Biomedical Applications

by Marc Thiriet

The first Workshop on Multidisciplinary Simulation in Biomedical Applications was held in cooperation with ERCIM Working Group "IT and Mathematics applied to Interventional Medicine" (IM2IM) at Berufsakademie (University for Cooperative Education), Mosbach, Germany from 21-22 November 2006.

The focus of the workshop was on multiphysics simulation, ie the use of different solvers, particularly coupled simulation methods mainly using MpCCI (Mesh-based parallel Code Coupling Interface), developed by the Fraunhofer SCAI team. MpCCI is the common feature of most of the work presented during the workshop:

- Gerhard Link, Sensor Technology, University of Erlangen, gave a talk on substitute voice by the pharyngeal-esophageal segment (PES) and the effect of geometry on voice quality.
- Christoph Müller, CADFEM GmbH - Ansys, presented a "patient-individual dental implant simulation", after briefly giving other simulation examples (spine column, bone implant planning, computation of muscle forces during body motions, gait analysis, surgery planning such as jaw remodelling, and forced extension of palate).
- JG Schmidt, NEC Research Lab, St Augustin, spoke on "integrated biomedical informatics for the management of cerebral aneurisms". The multi-centre study involving many academic and industrial partners and supported by the EU is aimed at combining data (genetics, morphodynamics, etc.) for rupture risk assessment.
- Andreas Burblies, Fraunhofer IFAM, Bremen, spoke on "material design for medical applications", the main application being the orthopedic field. The common goals are weight reduction and structure optimisation using hybrid multiphase structures.
- Johannes Meenen, Abaqus GmbH, gave a presentation on stenting and drug delivery from coated stents. Abaqus contains a standard implicit solver, suitable for slow mass transfer and also used to solve fluid-structure interaction, and an explicit dynamic solver, used when discontinuities occur (contact etc), especially for stent expansion and catheter retraction.
- Ralf Kröger, Ansys/Fluent, talked on microfluidics, focusing on neutral hydrodynamics (as opposed to electrohydrodynamics). Illustrations on surface tension effects were given, as well as on droplet production.
- Friedhelm Schönfeld, Institut für Microtechnik, Mainz, talked on biomicrofluidics (surface-tension-driven liquid flows which involve new physics, rarefied gas dynamics, etc) and the development of a modular chip-based lab platform.
- Jürgen Berieter-Hahn, Institut für Zellbiologie und Neurowissenschaft, Center of BioMedical Engineering (CBME), Frankfurt am Main, introduced the CBME. Decubitus, aneurisms and stent design are the main targets of research in CBME.
- Klaus Wolf, Fraunhofer SCAI, focused his talk on code coupling, using the neutral open interface MpCCI tool. It treats FSI, aeroacoustics, magnetohydrodynamics etc.
- Marc Thiriet, INRIA, presented numerical results in image-based models of physiological flows. He demonstrated how helpful computational results can be in medical practice. He also pointed out the main drawbacks of the methodology, leading to qualification rather than quantification. He then talked about modelling strategies, and the ancient version of 'reductionism' leading to modern 'integrative approaches' of complex physiological systems.

A half-day session meeting was devoted to the 7FP, in order to define a proposal skeleton for ICT second calls in 2007.

Please contact:
Marc Thiriet
INRIA, France
ERCIM IM2IM Working Group coordinator
E-mail: marc.thiriet@ercim.org

New Projects managed by ERCIM

Two new projects, VITALAS and EchoGRID, both administrated by ERCIM and partly funded by the European Commission commenced in January 2007.

VITALAS (Video & image Indexing and Retrieval in the Large Scale) is an 'Integrated Project' that will deliver a prototype system for accessing professional multimedia archives. The challenges addressed by VITALAS include cross-media indexing and retrieval, large scale search techniques, and visualisation and context adapting. The project consortium consists of eleven partners from research organisations and industry, including the ERCIM members CWI, Fraunhofer-Gesellschaft and INRIA. The project has an overall budget of €8 million for three years.

EchoGRID (European and Chinese Cooperation on Grid) is a so-called 'Specific Support Action' to foster collaboration in Grid research and technologies by defining a short-, mid- and long-term vision in the field. A set of thematic workshops on strategic topics defined by a group of experts will be organised in the frame of this projects. A research exchange and fellowship programme will also be set up to encourage the exchange of personnel across countries and institutions.

Links:
http://www.ercim.org/activity/projects/

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

W3C Celebrates Ten Years with Style

Ten years ago, on 17 December 1996, W3C published the first standard for style on the Web: Cascading Style Sheets (CSS), level 1. CSS Web designers have since enjoyed fine-grain control of page appearance (fonts, colors, layout, margins, etc.) and easier page design and maintenance. CSS can also help make pages more adaptable to more users, including users with mobile devices and some users with disabilities.

"The design community has confirmed that using CSS, and separating markup from presentation, promotes beauty while making it easier and less expensive to build sites, " said Bert

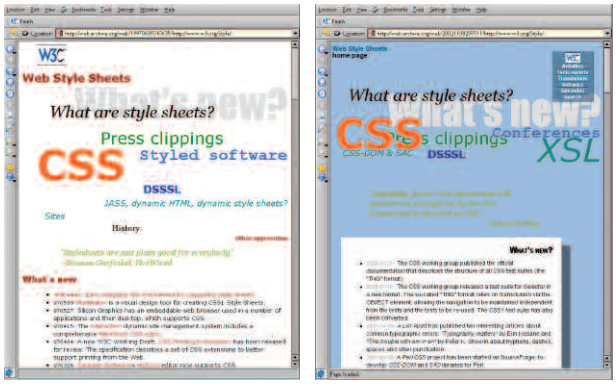

The W3C Style home page around 1997 and around 2002.

Bos, W3C Style Activity Lead and one of the original co-authors of the specification that became CSS level 1, published on 17 December 1996.

To celebrate this tenth anniversary, W3C invites developers to propose their favorite CSS designs for the CSS10 Gallery. Bert Bos and Håkon Lie, the other original co-author of CSS, will select designs for the gallery based on originality, utility, and aesthetics. Proposals are welcome until December 2007; we expect to add to the gallery at least once monthly. The proposed designs are to be sent to css10@w3.org.

In addition, thanks to the efforts of users, developers, and translators, W3C has released a new version of the CSS validator in time for CSS10.

CSS Separates Markup from Presentation, Benefitting Designers and Users Alike

CSS success derives from its numerous benefits to designers. The first benefit is the rich feature set. Using a simple declarative style, designers can set positioning, margins and alignment, layering, colors, text styling, list numbering, and much more. Furthermore, writing direction, font styles, and other conventions differ from one written language to another. CSS supports an increasing number of different typographic traditions and has made significant progress toward being able to display multilingual documents.

The second benefit is reuse. Style sheets can be shared by multiple pages, making it easy to update an entire site by changing a single line of CSS. Because style sheets can be cached, this can mean improved performance as well.

CSS promotes accessibility in a number of ways, without compromising design. Separating markup from style enables accessibility agents to convey information according to the needs of users with disabilities. The CSS design strikes a good balance between author and user needs, enabling users to make use of more pages. Style sheets also reduce dependency on using HTML tables for layout, which can be a barrier to some users with disabilities using assistive technologies such as screen readers.

A related CSS benefit is easier cross-media publishing; the same document may be viewed with different devices (from large color monitors to mobile phones to printers) simply by applying the appropriate style sheet. Software can choose the most appropriate style sheet automatically (as suggested by the style sheet author), and allow the user to choose from among available style sheets to meet that individual's needs.

CSS is commonly used to style HTML and can also be used with XML documents as a complement to W3C's XSL.

CSS3 Targets Multimedia, Multimodal

CSS has various levels and profiles. In general, desktop browsers implement level 1, 2 or 3. Other programs implement the appropriate profile for their platform, whether mobile phone, PDA, television, printer, speech synthesizer, or other device.

CSS level 1 defines properties for fonts, margins, colors, and other tools for style that are common to nearly all profiles of CSS. An early example of CSS use is the original CSS gallery (written when Microsoft's Internet Explorer 3 added CSS support). W3C has compiled more CSS history as part of this celebration.

CSS level 2 revision 1 (CSS 2.1) includes all of CSS level 1 and adds absolutely positioned elements, automatic numbering, page breaks, right to left text and other features.

CSS level 3 (CSS3), still in development, promises more power features at the same time it will make CSS easier to implement and use. CSS3 includes all of level 2 and adds new selectors, rich hypertext, more powerful borders and backgrounds, vertical text, user interaction (e.g., styling of XForms), speech, rendering on multimedia devices, and more; see the CSS Working Group charter for details.

Links:
CSS10: http://www.w3.org/Style/CSS10/
CSS10 Gallery:
http://www.w3.org/Style/CSS10/reactions.html
W3C Style Activity: http://www.w3.org/Style
CSS Validator: http://jigsaw.w3.org/css-validator/

W3C Workshop on Web of Services for Enterprise Computing

W3C is holding a Workshop hosted by MITRE, 27-28 February 2007 in Bedford, MA, USA, to discuss how to facilitate the processing of business transactions and interactions with systems that pre-date the Web, and to address the need to interconnect intranet and/or extranet services using Web technologies.

The Workshop is free and open to both W3C Members and non-members. Space is limited and participation may be limited by organisation. Interested parties are required to submit position papers no later than 10 January 2007.

Link:
CfP: http://www.w3.org/2006/10/wos-ec-cfp.html

Third W3C Workshop on Internationalising the Speech Synthesis Markup Language (SSML)

W3C held its third Workshop on Internationalizing the Speech Synthesis Markup Language (SSML). The Workshop was held at International Institute of Information Technology (IIIT) in Hyderabad, India, jointly hosted by Bhrigus Software and IIIT. As a result, various requirements on SSML extensions for Asian, Eastern European, and Middle Eastern speech synthesis were identified and prioritized.

Link:
CfP: http://www.w3.org/2006/10/SSML/cfp.html

Markup Validator, Link Checker Updated

W3C has updated its W3C Markup Validation Service and Link Checker with bug fixes, documentation and usability improvements, and a new Validator API for developers. Along with W3C's other Web Quality Tools, the Markup Validator and Link Checker are developed as open source software with the participation of volunteers and support of a large community, and are among W3C's most popular and useful resources.

Links:
Validation service: http://validator.w3.org/
Link checker: http://validator.w3.org/checklink
W3C Web Quality Tools: http://www.w3.org/QA/Tools/

W3C10 Asia: Tenth Anniversary of W3C in Asia

W3C10 Asia is the celebration of the tenth anniversary of the founding of W3C's Asian host at Keio University in Japan in September 1996. A public celebration was held on 28 November in Tokyo, Japan.

The one-day W3C10 Asia symposium consisted of three panel discussions with luminaries from Japan and Asia whose contributions have played a key role in promoting the Web, ensuring that the Web is open and accessible to everyone, and helping the Web to reach its full potential in Asia. The program included "Role of W3C at Keio - From Foundations to the Future," "How Japanese Industry Works with Web Standards," "How Asia Will Influence the Future Web," discussions.

Celebrating W3C10 Asia with Sake Breaking Ceremony

The symposium closed with reflections and projections in a dialog with Dr. Jun Murai, Vice-President of Keio University, and Tim Berners-Lee, KBE, inventor of the World Wide Web and founder of W3C.

Link:
W3C10 Asia: http://www.w3.org/2006/11/W3C10/

Latest W3C Recommendation

Extensible Stylesheet Language (XSL) Version 1.1
5 December 2006, Anders Berglund

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

The European Strategy Forum for Research Infrastructures

by John Wood

In October 2006 the European Strategy Forum for Research Infrastructures (ESFRI) produced its first roadmap. This identifies new Research Infrastructure (RI) of pan-European interest, which will help develop the European Research Area (ERA) over the next ten to twenty years. This has now been handed over for consideration by the Commissioner for Research and the science ministers of the EU member states.

In my letter to them as chair of ESFRI, I state that "Some are evolving projects while others need firm financial commitments from the start. The role of ESFRI in fostering incubation and stimulation will be exercised to bring as many of these projects to a point where decisions by ministers are possible. This requires in the first place, discussions and decisions at national level in particular as regards the lead role one country or several countries may wish to take for certain projects." This point underlines the key role of member states in taking the initiative. Many external observers have cherished the mistaken belief that the roadmap is a list of projects that the EU will fund. The European Commission wishes to aid research and is looking at ways of partly funding the costs of project set-up or development. However, the vast majority of funding must come from member states where they feel a particular project should be of high priority. The process of creating this document is a remarkable achievement, and demonstrates the truly international perspective of most researchers. Nevertheless, a number of issues had to be put to one side in order to deliver a document that would be credible and on time.

One that needs further elaboration is how e-science and computing infrastructure fit with these specific distinct European facilities. Before looking at the implications of this, it is worth reflecting on what ESFRI was set up to do, what it does not do, how it functions and how the roadmap was put together. This will highlight clearly the sort of approach that needs to be integrated with the actions of e-IRG (EC e-Infrastructure Reflection Group) and its future aspirations in developing e-science and supercomputing capability. In February 2002, a high-level expert group comprising representatives from all member states issued a report recommending the creation of ESFRI. Following this, M. Busquin wrote to science ministers of the fifteen member states asking for national nominees to attend the Forum and offering operational support from the Commission.

The first meeting took place in April 2002. Delegates are appointed for a two-year renewable term and must have a direct link to funding bodies and decision makers in their respective countries. In addition, there are observers from potential accession and associate countries, representatives from a number of European organisations, and invited speakers. This means that Forum meetings can have up to a hundred attendees. The Forum's principal objectives are to support a coherent approach to developing policy on research infrastructure in Europe, and to facilitate multi-lateral discussions leading to their better use and development.. The Forum acts as an informal body on issues raised by one or more delegations. It does not make decisions on priorities, funding or the ways in which research infrastructure is subsequently managed. ESFRI gives national authorities the opportunity to be informed of and to explore initiatives concerning the development or upgrading of research infrastructure in areas of European significance. In this way it takes on the role of an incubator. ESFRI is also concerned with other aspects of research development. For instance, it looks at models that have and could be used for pan-European funding and management, at suitable metrics for assessing the value-for-money and socio-economic aspects of such investments, and at how transnational opportunities for young researchers can best be encouraged. It became apparent in 2004 that a number of countries – European and non-European – were developing long-term (ten to twenty year) roadmaps for research infrastructure. These showed policy makers and funding bodies the areas in which specific research communities aspired to develop, and the approximate funding required. As a result, during the Dutch presidency of the EU, ESFRI was asked to undertake a similar exercise, building on national strategies where possible.

Key challenges facing Europe in the XXI century: the contribution of Research Infrastructures.

This was a formidable exercise and has dominated the activities of ESFRI until publication. It was decided that all areas of research should be covered including

social sciences and humanities. There were discussions on how the exercise should be undertaken. It was decided that three working groups should be set up to investigate specific areas along traditional discipline lines. Many argued for themes such as the environment, pointing out that these cut across disciplines. It was also noted that several areas of research infrastructure supported more than one discipline.

The modus operandi and criteria are fully outlined in the published roadmap. A number of expert groups reporting to the working groups were formed and over a thousand researchers took part in the exercise of analysing proposals. Around two hundred were involved in the peer review of individual projects that were proposed by ESFRI delegates or recognized European organisations. I was delighted by how quickly there emerged a consensus on the key projects that were considered mature enough to be incorporated onto the roadmap, and the final list was unanimously approved by the ESFRI delegates.

For ease of use by policy makers, the areas of infrastructure were subsequently listed under key themes (energy, materials science etc). However, these are somewhat arbitrary, and a series of commentaries is therefore included to elaborate on how each area will contribute to the important challenges facing society. During the construction of the roadmap, discussions started with e-IRG and DG-INFSO on the integration of their developing ideas. Initially it was felt that there might be competition, but with good chairmanship of e-IRG and an openness of ESFRI to get things right, e-IRG delegates attended the Physical Science and Engineering working group. The chair of e-IRG is now permanently invited to the ESFRI meetings.

There is no doubt that the role of e-infrastructure is all pervasive, especially as we move towards remotely driven experiments. However the range of issues is much larger. First, there is the development of massive European databases following population studies, future developments in biosciences and so on. This raises several points. The first surrounds archiving, curation and access to these data. The second relates to verification or provenance, which is a more pressing issue than many realise. The question of how to prove a set of data/metadata is true and has not been corrupted needs to be actively addressed. I often talk about the situation where structural biological data could be corrupted and subsequently used in the design of drugs for field testing. Who will responsible for the millions of euros spent? What about the cost to human lives and to the reputation of the company subsequently employing the results? We do not need to look too far into the recent past to see that data has been published that is untrue. How do we check, who is the final arbiter? As the volume of published data increases, the seriousness of the situation increases. A third point is the linking of simulations with the real-time control of experiments. Some of the infrastructure areas will result in the data flows of millions of pixels on near-femtosecond timescales. As the data comes in it will be necessary to alter the experimental conditions in similar time regimes in order to gain maximum output from a very expensive plant. Computer simulations are required that take in the real data, adapt the models and then issue instructions to the equipment on a time scale of which no human experimenter can conceive. The necessary operating systems, codes and feedback loops need to be developed to ensure that this can take place.

Then we come to the hierarchy of computers needed to support such research. How many supercomputers does the EU need? Here we face the dilemma of juxtaposing the need for capacity (ie a fast number cruncher) with the need for machines designed for specific models. Prof. Tetsuya Sato, the Director-General of the Earth Simulator Center in Japan, argues that both are needed within physical proximity of a few hundred metres. The dedicated machines develop the models which are then handed over to the number crunchers. Is this the right way for Europe? The ESFRI roadmap highlights these issues, but as yet has not resolved them other than to say a hierarchy is needed. The exact words are: "The facility: A European strategic approach to high-performance computing concentrating the resources in a limited number of world top-tier centres in an overall infrastructure connected with associated national, regional and local centres forming a scientific computing network to utilise the top-level machines. This overall architecture will respond both to the Capability (high performance) and Capacity Computing (high throughput) needs. Different machine architectures will fulfil the requirements of different scientific domains and applications. This can be represented as a pyramid where local centres would constitute the base of the pyramid, national and regional centres would constitute the middle layer and the high-end HPC centres would constitute the top." The question that must now be answered is how best to develop and implement these ideas in such a way as to deliver meaningful solutions to European society in the future.

ESFRI is very proud to present the first European roadmap for new, large-scale Research Infrastructures, based on an international peer-review.

Link:
http://www.cordis.europa.eu/esfri/

Please contact:
John Wood, CCLRC, Chair of ESFRI
E-mail: J.V.Wood@cclrc.ac.uk

John Wood.

Europe: The Place to Be!

by Peter Nijkamp

EuroHORCs (European Heads of Research Councils) is a platform for principals (chairs, presidents, CEOs) of national research organisations and funding bodies to meet to discuss common issues and develop new initiatives. The organisation bears an enormous responsibility to the research landscape in Europe, having a joint annual budget exceeding 20 billion euros - far higher than the EU Framework Programme (FP) budget. Thus, the organisations represented in EuroHORCs are essential players in both national research systems and in the European Research Area (ERA).

While in the past Europe has developed a patchwork of nationally-based research funding systems, it is increasingly recognizing the need to sharpen its focus in R&D, in order to cope with emerging global scientific challenges. Europe cannot afford to lose its research talent in a global science market and must therefore take effective action in order to become a favourite research region for promising talent.

The European Response

There is a clear need for a more coherent and effective joint research policy in Europe. Recent decades have already witnessed a gradual transition towards greater cooperation, mainly in the context of the various FPs. However, a rigorous attempt to achieve coordination between cutting-edge European research programmes has so far been missing, and new initiatives are necessary. The linking and integration of national research programmes funded by the EU and by national funds is a potentially powerful mechanism for revitalizing the ERA, provided such a mechanism aids scientific performance and administrative efficiency. This requires that R&D organisations maintain a level of independence; in addition, the coordination of national research programmes should be transparent and peer-reviewed, and demonstrate a clear value to science. Only if these conditions are met will it be possible to attract the best scientists to Europe, as is witnessed in CERN or EMBL.

The European Partnerships

There are many advantages to a partnership among R&D stakeholders in Europe. For instance, it will be able to promote:

- the optimal nurture and growth of excellence
- flexible career paths for (young) promising researchers from all over the world
- efficient coordination of research programming and funding in Europe, and in other countries
- cross-border cooperation and joint publications among European scientists
- funding of projects and programmes with European impact and strategic intent (added value)
- scale and scope at the European level, thus strengthening the role of Europe in the global arena
- learning from (and making use of) each other´s funding systems (access)
- the creation and strengthening of critical mass for the protection of the viability and vitality of 'small disciplines'
- the setting up of a European system for review, benchmarking and best practice
- better integration between (formerly) Eastern and Western Europe.

Common multilateral R&D activities are also necessary for the revival of the ERA from a strategic long-range European perspective. Even with a European Research Council (ERC) established in FP7, frontier, innovative and high-risk research will for the next decade tend to be relatively under-represented in EU funding. In addition, there is still too little funding for bottom-up, international cooperation in frontier research, since FPs are traditionally aimed at strategic research in areas of societal or industrial priority. Europe must therefore seek a balanced research policy.

The EuroHORCs Initiatives

EuroHORCs has developed a rich portfolio of research initiatives to serve the needs of the scientific community in Europe. These include:

- An ongoing dialogue with the European Commission on the modalities and tasks of the ERC. This permanent interaction – in a very constructive atmosphere – has been of decisive importance in shaping the ERC.
- The creation of the EURYI (European Young Investigator) scheme, which served to scout and facilitate top-level young researchers. This scheme has an excellent reputation and has laid the foundation for the Starting Independent Researcher grand scheme of the ERC.
- The linking of nationally funded research programmes in Europe (such as Eurocores and Eranets) in order to reduce fragmentation. This scheme has laid the cornerstones for further cooperation among research funding agencies in Europe. Further, future perspectives on new forms of cooperation have been developed in close collaboration with the European Commission.
- Mutual service provision by European research councils and institutes in a variety of areas. Examples include joint initiatives on research infrastructures, stimulation of international mobility through Money-Follows-Researcher schemes, development of common ideas on IPR, linking of databases of referees, and development of common standards for peer review systems. Many EuroHORCs initiatives are jointly undertaken with the operating arm of EuroHORCs, the European Science Foundation (ESF).

In conclusion, EuroHORCs has a booming action agenda, and it is foreseeable that its role in the ERA will become even more critical in years to come.

Link:
http://www.eurohorcs.org/

Please contact:
Peter Nijkamp, President NWO,
President EuroHORCs
E-mail: eurohorcs@nwo.nl

Picture: NWO / Ivar Pel.

Peter Nijkamp

CERIF - the Common European Research Information Format

by Keith G Jeffery

Information technology has been used to support parts of the research process for many years. Research funding organisations have systems to manage project proposals and grants awarded. Universities have systems to keep track of research projects. Laboratory notebook systems have been developed, some with a workflow component. Portals to access research publications and research datasets and software have been developed. However, these systems tend to be stand-alone, heterogeneous, incompatible with each other and of varying stages of maturity.

The Common European Research Information Format (CERIF) was developed with the support of the European Commission (EC) in two major phases: from 1987-1990 and 1997-1999. It is a standard; technically it is a recommendation by the European Union to its member states. Since 2002, care and custody of CERIF has been handed by the European Commission to euroCRIS, a not-for-profit organisation dedicated to the promotion of CRIS (Current Research Information Systems).

The rationale behind CERIF is as follows. Most nation-states maintain publicly supported research programmes. It is realized that public sponsorship of research and development leads to wealth creation and improvement in the quality of life. Because public funding is involved, it is necessary for there to be appropriate governance, and also for the information to be made available to the public. Broadly, each nation state has a similar research process, which involves strategic planning, programme announcement, a call for proposals, proposal evaluation and awarding, project result monitoring, and project result exploitation.

However, research is international. A research project in country A is likely to be based on previous research in several other countries. Many research projects are now transnational: well-known examples include the human genome and climate change but there are many others, especially where expensive infrastructure is utilized such as particle physics or space science. Furthermore, knowledge of the research activity in country A may influence the research strategy – including priorities and resources provided – of country B.

Thus there is a need to share research information across countries, or even between different funding agencies in the same country. The information is used by researchers (to find partners, to track competitors, to form collaborations); research managers (to assess performance and research outputs and to find reviewers for research proposals); research strategists (to decide on priorities and resourcing compared with other countries); publication editors (to find reviewers and potential authors); intermediaries/brokers (to find research products and ideas that can - with knowledge transfer - be transformed to wealth-creation); the media (to communicate the results of R&D in a socio-economic context) and the general public (for interest).

Within Europe the importance of sharing information has long been recognised. In the early eighties, the heads of some national research funding organisations initiated a project named IDEAS to investigate linking databases of research information. The follow-on project EXIRPTS (1987-1989) extended this – based on the G7 countries - to include USA and Japan. Both projects successfully demonstrated: (a) resolution of schema differences between heterogeneous distributed databases; (b) translation and distribution of queries as subqueries to the different target systems; and (c) integration of the results of the distributed subqueries and presentation back to the queryier. The technology utilized a specially developed protocol, run over email and file transfer. It should be noted that this was well before the emergence of the World Wide Web.

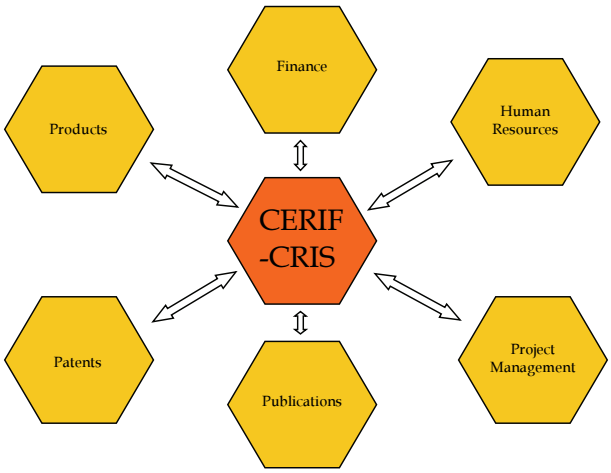

CERIF CRIS can be used to relate the research output product to any CERIF entity such as project, person or organisational unit or publications.

It is from this work that CERIF emerged, first as a simple standard (not unlike a library catalogue card or the present Dublin Core Metadata Standard). It was intended as a data exchange format, and was based on records describing projects, with persons and organisational units as attributes. However, it was soon realized that in practice this CERIF91 standard was inadequate: it was too rigid in format,

did not handle repeating groups of information, did not support multiple languages and character sets, and did not represent in a sufficiently rich way the universe of interest.

In response, a new group of experts was convened and CERIF2000 was generated. The latest version (CERIF2006) is described in detail at http://www.eurocris.org/, with its essential features being:

- it has the concept of objects or entities such as project, person, or organisational unit with attributes
- it supports n:m relationships between them (and recursively on any of them) using 'linking relations', thus providing rich semantics including roles and time
- it is fully internationalized in language and character set
- it is extensible without prejudicing the core datamodel, thus providing guaranteed interoperability at least at the core level while not precluding even richer intercommunication. It is designed both for data exchange (data file transfer) and for heterogeneous distributed query/result environments.

CERIF is now in use in numerous organisations. One commercial company (uniCRIS AG, http://www.unicris.com/) offers CERIF-compatible systems and three others have stated that they will evolve their system offerings to do so. The European Science Foundation has specified CERIF in its new IT system and CORDIS, the EU R&D Web site, is evaluating how best to be CERIF-compatible. Several ERA-NETs funded by the EC use CERIF to allow interoperation of information on research in each member state. The projects IST-World and CISTRANA, both aimed at providing an overview catalogue of European research, use CERIF.

Within euroCRIS, an active task group is developing CERIF and provides advice on how best to use it. To participate, anyone interested is requested to contact EuroCRIS.

Links:
EuroCRIS: http://www.eurocris.org
Keith G Jeffery, Anne Asserson: 'Supporting the Research Process with a CRIS': http://epubs.cclrc.ac.uk/work-details?w=35466

Keith G Jeffery, Anne Asserson: 'CRIS Central Relating Information System': http://epubs.cclrc.ac.uk/work-details?w=35484

Please contact:
Keith Jeffery
CCLRC, UK
President of euroCRIS
President of ERCIM
E-mail: kgj@rl.ac.uk

Keith Jeffery

Managing Research and Knowledge Based Organisations

by Günter Koch

As competition for research funds becomes tougher, professional management of science and R&D plays an increasingly crucial role.

In the past, the typical profile of leaders in research was that they had to qualify as scientists following the specific career rules of the community, and perhaps also had talent or intuition to effectively manage intellectual teams or organisations. Good management means a) providing conditions that enable employees to fully develop and demonstrate their potential, and b) preparing decisions and having them implemented. In recent years it has become evident that academic institutions, specifically in the Anglo-American world, can benefit from management potential developed in business environments. The constitutions of modern universities (eg as expressed in the Austrian University Law of 2002) even explicitly foresee the mandate for

Figure 1: The RICARDIS Report, availble at http://ec.europa.eu/invest-in-research/pdf/download_en/2006-2977_web1.pdf

employees with backgrounds other than science or research (ie business and management) to take leading roles.

Nevertheless, there are some remarkable differences between companies and educational or research organisations. The most evident is that companies are usually devoted to making profits in monetary terms, whereas the 'currency' of the science community is the scientific publication. Digging deeper into this subject, it turns out that there are many more indicators than simply profit, by which an educational or research organisation should be measured. These include the satisfaction and intellectual productivity of employees, the culture of the organisation, involvement in networks, the

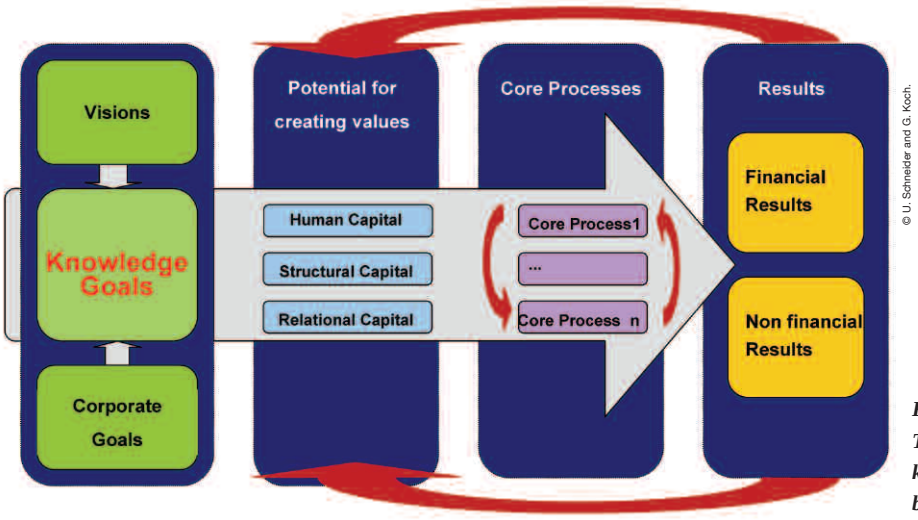

Figure 2:
The "Standard Model" of a knowledge and research based organisation.

existence of quality management, and the 'production' of academics for a society. All these metrics contribute to the quality, and in the end, to the excellence and profile of a research institution. This is usually reflected by rankings (eg of universities), which have become popular and influential over the last decade.

In the 1990s, the idea that an organisation may be characterized not only by economic performance but also by its potential in knowledge and know-how was demonstrated by the definition of the role of Chief Knowledge Officer (CKO), first installed at the Scandinavian insurance company Skandia in 1994. This idea, combined with quality models taken over from Software Process Quality & Maturity as reflected in standards such as ISO 15504/SPICE, was taken up by R&D managers everywhere. The trend was also observed in a majority of large organisations in applied research, such as the ERCIM members VTT, Fraunhofer, SINTEF and ARC, all of which are assembled in the European Association of Research and Technology Organisations (EARTO).

As requested by the European Commission, they constituted a 'high-level expert group', which recently came up with the RICARDIS report (see Figure 1). In substance, this report is a compendium of methods for managing research- and knowledge-based organisations, both public and private. (As the case of Austria proves, even universities have accepted such methods of professional knowledge management and now issue an annual 'intellectual capital report').
Figure 2 illustrates the important aspects of management in an R&D organisation. This reference model combines all the relevant dimensions of a knowledge organisation in one framework and, at least in the German-speaking countries, has become a kind of standard. In Germany, the government is promoting intellectual capital reporting in order to motivate knowledge-based companies to advertise their potential and to make knowledge management a strategic objective.

The ultimate goals of such a 'Knowledge = R&D Management' framework are to raise the quality of R&D management, and to put R&D organisations in a position to become 'benchmarkable' by their own criteria, rather than by rankings imposed on them from the outside.

It is estimated that already around 500 organisations in Europe apply intellectual capital reporting and the associated management methods. There are indications that the EC will soon extend the existing EU standard in financial reporting (IFRS) through a directive asking for additional information on the ability of an organisation to produce, aggregate and exploit knowledge generated from R&D.

In 2006, ERCIM's Board of Directors decided for the first time to provide an intellectual capital report as an addendum to its Annual Report, to be issued in 2007.

Link:
RICARDIS report:
http://ec.europa.eu/invest-in-research/pdf/download_en/2006-2977_web1.pdf

Please contact:
Günter Koch, execupery @ TechGate / AARIT, Austria
E-mail: koch@execupery.com

Günter Koch

Introduction to the Special Theme

Traffic Planning and Logistics

by Jo van Nunen

The ambition of the European Union is to take a leading position in the world economy. Crucial to the business processes that can bring such prosperity to European citizens are sustainable transport and logistics systems. Successfully developing such systems requires a balance between the three 'Ps' of sustainability – Planet, People and Profit. This is not easy, since transportation has a number of negative consequences, including congestion and undesirable emissions. In this issue of ERCIM News, scientists in the field of traffic management, transport and logistics demonstrate how mathematical modelling can help to find the right balance between the 3 'Ps' while creating a sustainable future for the transport and logistics sector. The role of information technology is also discussed, since it is increasingly important in providing the infrastructure for monitoring traffic and logistics systems, for handling collected information and for providing the algorithms and methods necessary to analyse the large amounts of data that are now available.

First, consider the possibilities of using different devices for sensing logistic flows. Passive RFID-tags (radio frequency identification tags) mounted on cargo allow the flow of goods to be measured every time the cargo passes an RFID reading device. By using active RFID tags combined with GPRS devices, cargo flows can even be followed continuously. The information stored in the tag typically includes the type of product, the producer, the planned destination, the consignee and so on. During the transport process this information can be used to optimize the flow of goods, for security controls by customs at transshipment points and for sorting goods with similar destinations into combined shipments. This allows safe and secure flows to be created within designated trade lanes, and for the flows to be designed in such a way as to minimize environmental impacts.

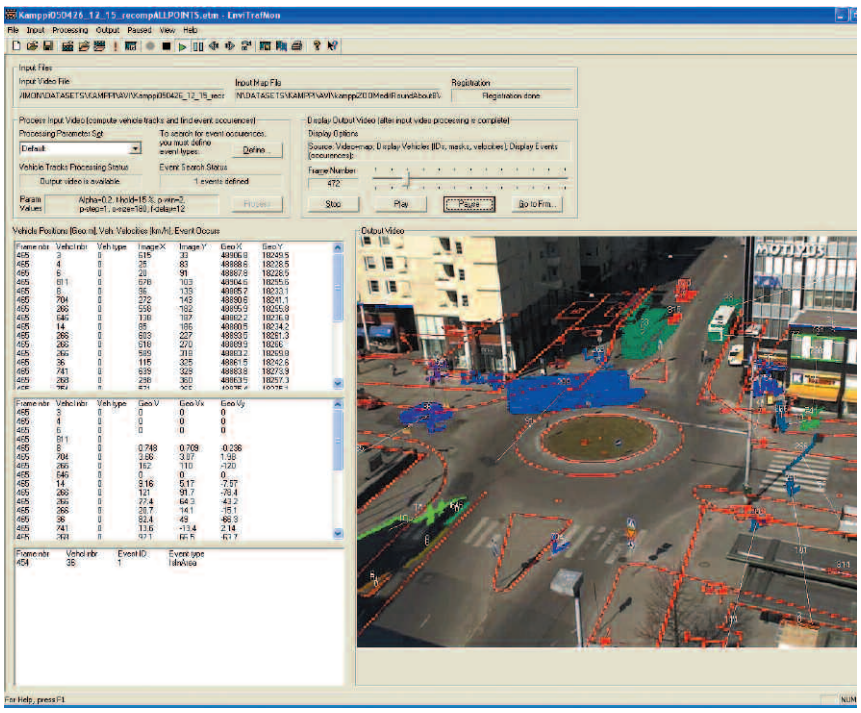

Screenshot of TRAFMON, a tool for interactive traffic monitoring (see article on page 51). Information technology is becoming evermore crucial in providing the infrastructure for monitoring traffic and logistics.

If the cargo can be monitored on a real-time basis, additional opportunities to improve logistics processes become possible. Providing subsequent parties in the supply chain with advance information on the progress of a logistical process enables those parties to optimally plan their operations. In this issue several examples are discussed, varying from tracking and tracing of information to collaborative navigation solutions. If we extend the definition of logistics to include the production phase at the front end of the chain coupled with the treatment of the products after they are used, we enter the field of reverse logistics and closed-loop supply chains. Research in this area focuses on the possibilities of reusing, refurbishing, remanufacturing and recycling products and materials in a sustainable way. European research is at the forefront of this field and offers many opportunities for new economic activity. Reverse logistics research is based on the concept that “there is no waste, all products and materials are resources”, something that fits well with the European value system.

Closed-loop supply networks, in which products and carriers can be monitored over their complete lifetime leads to the concept of installed base management. Cars, copiers, computers and machines, as well as objects such as

Photo: Schenker AG.

Cargo flow can be followed continuously with an active RFID tag combined with GPRS devices.

tyres or lamps, can be monitored in real time. Processes such as repair and reuse can be triggered by the actual status of the installed product, which could even lead to a control shift "from tracking and tracing to sensing and pacing". As a consequence of this, work is needed on the phenomena of demand-driven and distributed control. Often virtual agent-based approaches can be used for this type of application. Within logistics research we are therefore seeing a shift from the solution of planning problems to online control approaches for which new models must be developed.

The problem of monitoring the flow of goods has many similarities to the monitoring of traffic and passenger flows. Personal mobile phones enable interactions between users of traffic infrastructure. Navigation systems with two-way communication, already common in truck transport, will be introduced for passenger transport. This will enable us to design intelligent transport systems that provide users with customized route selection choices in order to capture their preferred balance between Planet, People and Profit. Road-pricing and rewards for avoiding rush hour peaks can be dynamically incorporated in these customized systems. The models used in these cases incorporate multiple criteria in making decisions.

The navigation systems of the near future will use real-time information on road congestion based on mobile phone information from telecom operators. Using this information effectively requires that new methods be developed in order to forecast congestion, and research continues into advanced statistics-based methods for these types of applications. Further innovations in portable navigation systems will allow multi-modal transport solutions to be constructed, incorporating bike, bus, train and/or taxi in one trip and allowing users to select sustainable solutions to meet their needs.

At an aggregated level, the use of real-time information from all users of the infrastructure will lead to new forms of dynamic traffic management.

It is a challenge for us as researchers in this innovative field to contribute to these developments. In this issue of ERCIM News, you will read about a wide range of responses to this challenge.

Links:
http://www.euro-online.org/
http://www.rsm.nl/department1

Please contact:
Jo van Nunen
RSM Erasmus University Rotterdam
The Netherlands
E-mail: jnunen@rsm.nl

Mathematics for Railway Timetabling

by Leo Kroon

The Dutch railway timetable for 2007 is probably the only example of high-brow mathematics that is discussed by the whole population of the Netherlands. This timetable is partly based on mathematical optimisation models that were developed by CWI, Erasmus University, and others.

A new railway timetable was introduced in the Netherlands on December 10, 2006. This timetable is cyclic and involves both passenger and freight trains. One hour of the new timetable between the cities Gouda and Utrecht is shown in Figure 1. In this paper we describe several combinatorial optimisation models that were developed for generating cyclic timetables. These models played an indispensable role in the planning process of the new Dutch timetable.

Cyclic Timetable

Schrijver and Steenbeek (CWI) developed a model for generating the basic structure of a cyclic timetable. Their model assumes that the infrastructure and the line system are given, as well as the connections for passengers or rolling stock that have to be realized between certain lines at certain stations. Moreover, for all processes in the timetable (like running between stations, and headways between trains), minimum and/or maximum process times are defined. The model then looks for appropriate departure and arrival times for all trains at the corresponding stations.

The problem of generating the basic structure of a cyclic railway timetable can be described as a Periodic Event Scheduling Problem. It can be represented by a directed constraint graph, where the nodes are the departure and arrival times and the arcs are the processes between the departure and arrival times.

The Periodic Event Scheduling Problem can be solved by applying constraint propagation techniques. These techniques are adequate if the railway infrastructure is utilized intensively, as is the situation in the Netherlands. In that case, having fixed certain well-chosen parts of the timetable, the remaining solution space for the rest of the timetable is small, resulting in relatively short computation times.

Routing Trains through Stations

In the first timetabling step, the stations are considered more or less as black boxes: the details of the routes of the trains through the stations are considered only in an aggregated way. Therefore, in a second step, it must be checked in detail whether, given the arrival and departure times that were generated in the first step, the trains can be routed through the stations and allocated to platforms.

The problem of routing a set of trains through a station can be solved by first listing for each train the feasible routes through the station. Next, each combination of a train and a feasible route can be represented by a node in a graph. Two nodes are connected by an edge if they belong to the same train, or if the corresponding combinations of trains and routes are conflicting. The routing problem thus reduces to the problem of finding a maximum weighted node packing in this graph.

This weighted node packing problem can be solved by applying the commercial mathematical programming opti-

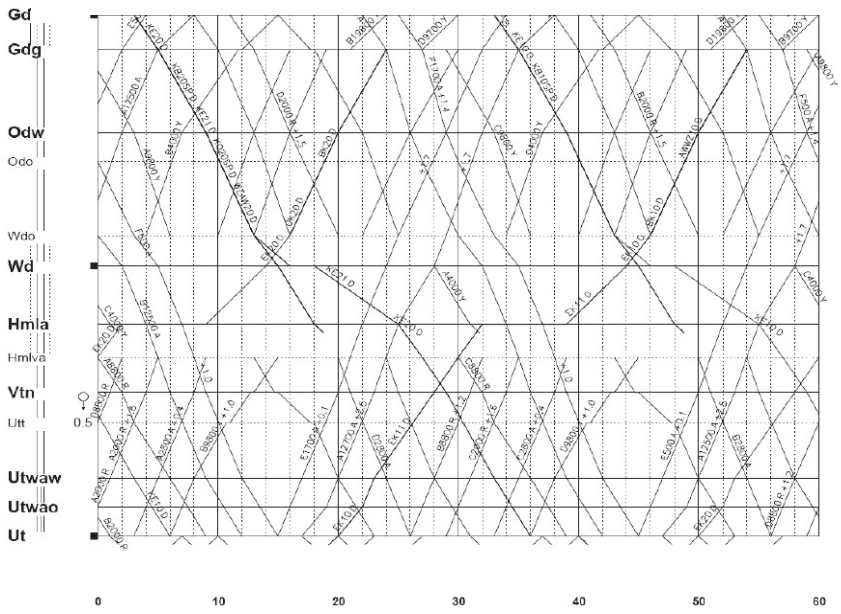

Figure 1: One hour of the new timetable between Gouda (Gd) and Utrecht (Ut).

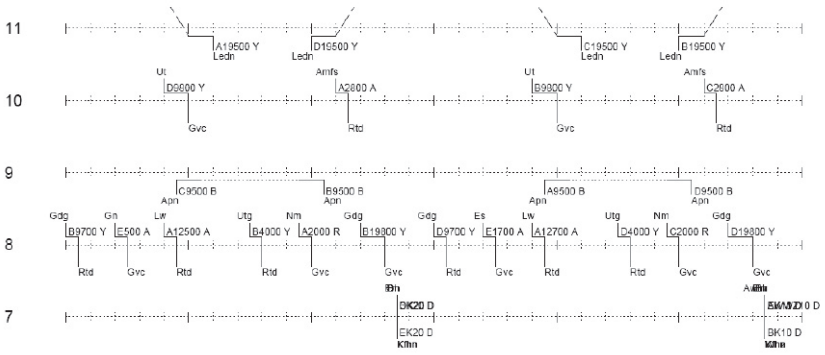

Figure 2: A platform allocation chart for tracks 7 to 11 in Gouda.

mizer CPLEX, after several pre-processing techniques have been applied for reducing the size of the graph. Figure 2 shows part of the platform allocation chart for Gouda station in the new timetable.

Robustness

In the first timetabling steps, it is assumed that the process times in the timetable are deterministic. However, the real-time process times in the operations are stochastic. The robustness of the timetable against small disturbances can be improved by optimally allocating time supplements and buffer times in the timetable.

This optimisation can be supported by a stochastic optimisation model that considers the processes both outside and inside the stations. The model modifies an initially given cyclic timetable and, at the same time, evaluates the modifications by simulating the trains in the timetable under stochastic disturbances. The aim is to minimize the average delay of the trains. This model can be considered as a symbiosis of an optimisation model and a simulation model.

Conclusions

The models described here were used intensively in the development of the Dutch timetable for 2007. The same holds for the multi-commodity flow model that was developed for planning the corresponding rolling stock circulations and for the set-covering model that was developed for generating the crew schedules.

Altogether, mathematics played an indispensable role in the generation of the new timetable, the rolling stock circulation and the crew schedules. The mathematical models allowed several scenarios to be generated, and explicit trade-offs to be made between conflicting optimisation criteria. This leads to a higher overall quality of the plans, as well as to a reduction of the lead-time of the planning process.

Currently, we are investigating whether these models can also be applied to support real-time control processes. That is, how can plans be adapted in the case of a real-time disruption of the railway system such that the passenger service remains as good as possible? Here the time needed for computing appropriate solutions is more important than near-optimality of the solutions: heuristic methods will be required.

Links:

http://www.ecopt.nl
http://www.cwi.nl/pna1
http://www.ns.nl

Please contact:
Leo Kroon
NS Reizigers, Utrecht and Erasmus University Rotterdam, the Netherlands
Tel: +31 30 2356658
or +31 10 4082421
E-mail: lkroon@rsm.nl

Railyard Shunting: A Challenge for Combinatorial Optimisation

by Per Kreuger and Martin Aronsson

Efficient railyard shunting is essential to cargo transporters in rail networks. In this article we explore methods of handling the situation where the capacity of the shunting yard is insufficient to handle all outgoing trains.

Sweden's largest rail cargo operator, Green Cargo, performs mainly three types of cargo transport: postal services consisting of fast point-to-point transports; system train services dedicated to large flows for particular customers; and wagon-load services. Railyard shunting occurs mainly for the third type.

In the wagon-load system, individual wagons are routed from point to point through a rail network generally consisting of five shunting yards of varying sizes throughout Sweden. Wagons are transported in trains that are assembled and disassembled at the shunting yards, each of which typically handles three or four incoming trains per hour. The flow of wagons through the network varies from day to day depending on demand.

Railyard Shunting

At the shunting yard, incoming trains are scheduled to arrive at fixed times but can use the entry group as a buffer and preparation area. Each incoming train is pushed over the barrier, where wagons are separated and roll down to the shunting group. Here switches are operated so as to distribute the wagons to a number of destination tracks. Figure 1 shows the layout of the medium-sized shunting yard.

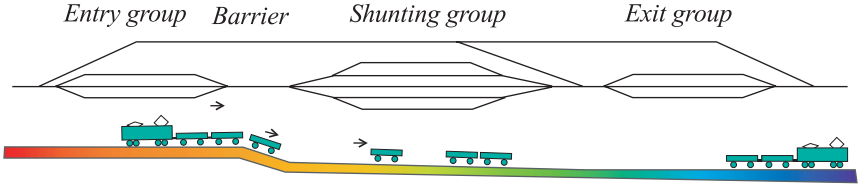

Figure 1: Schematic topography of a simple shunting yard.

There are several limited resources at the shunting yard:

- the entry group must be able to accommodate all incoming trains
- the barrier has a fixed capacity in terms of wagons per hour
- the shunting group limits the number of trains that can simultaneously be assembled.

Under 'normal' operation, a unique departing train is assembled at each destination track and departs at a predetermined time. In practice, this is not always feasible, since the capacity of the

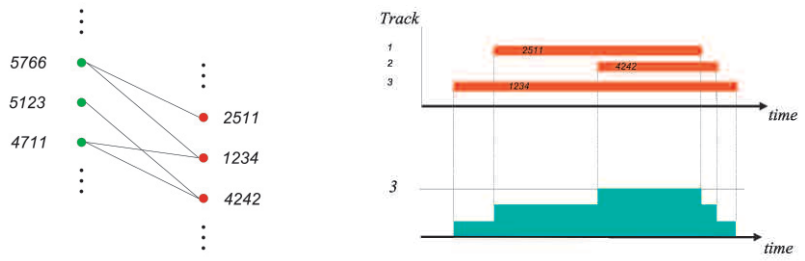

Figure 2: Flow and resource allocation for a case without temporary trains.

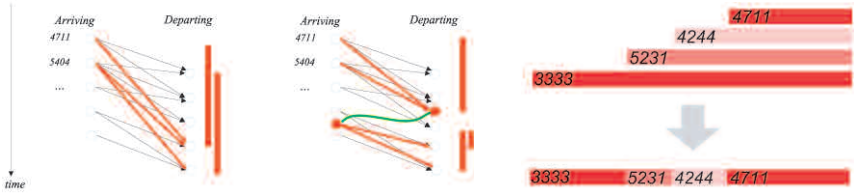

Figure 3: Flow and resource allocation for a case with temporary trains.

shunting group is in most cases too small for peak hour loads. To handle this, temporary trains are assembled and routed back to the entry group for re-shunting. This solution, however, leaves us with the following two considerations: which wagons should be combined into a temporary train, and when should this be routed back to the entry group and re-shunted?

Resource Considerations

We note that *without* temporary trains, the flow from incoming trains to outgoing trains is fixed. Figure 2 shows a small example of a typical flow and the corresponding Gantt chart and resource usage histogram.

With temporary trains, wagons destined for distinct departing trains may be routed to a single temporary train to reduce the capacity utilisation of shunting group. In order for this to be feasible, there must be sufficient time to the scheduled departure of each wagon to allow it to be re-shunted. The entry group and barrier must also be able to accommodate the corresponding additional tasks. Figure 3 shows the original flow, the modified flow with an additional temporary train and the resulting reduction in track resource utilisation.

A Tentative Solution

One possible model is to use a multi-commodity flow model as follows. Encode the destination train of each wagon as its commodity and allow incoming wagons to flow either directly to their destination or to a temporary train.

An outgoing train has potential inflows from each incoming train containing wagons destined for it and from all temporary trains. Temporary trains have individual flow conditions for each commodity. For each incoming train and commodity the sum of flows directly to the destination and to the temporary trains must be equal to that transported by the train. For each outgoing train the sum of flows from incoming and temporary trains must equal that transported by the train.

The scheduling constraints are augmented with tasks for each potential temporary train. Departure and duration for such tasks are constrained by the transferred commodities. Flow through temporary trains is restricted by the latest departure, which is in turn constrained by transferred commodities.

The model is implemented as a constraint program enforcing the multi-commodity flow conditions and scheduling constraints. In addition, the solver ensures that at any given time, at most one temporary train can be active.

Preliminary Results

Trial runs have been performed using real case data from a number of actual shunting sites in Sweden. For each such case the expected arrival and departure times were used. A small number of potential temporary trains with large arrival and departure time windows was introduced. The original flows in the case data were relaxed to allow each wagon to either flow directly to its destination or to a temporary train, which in turn has possible flows to each outgoing train.

For one temporary train and eight outgoing trains, correct solutions are computed in the order of seconds and for up to three temporary trains and on the order of forty outgoing trains, correct

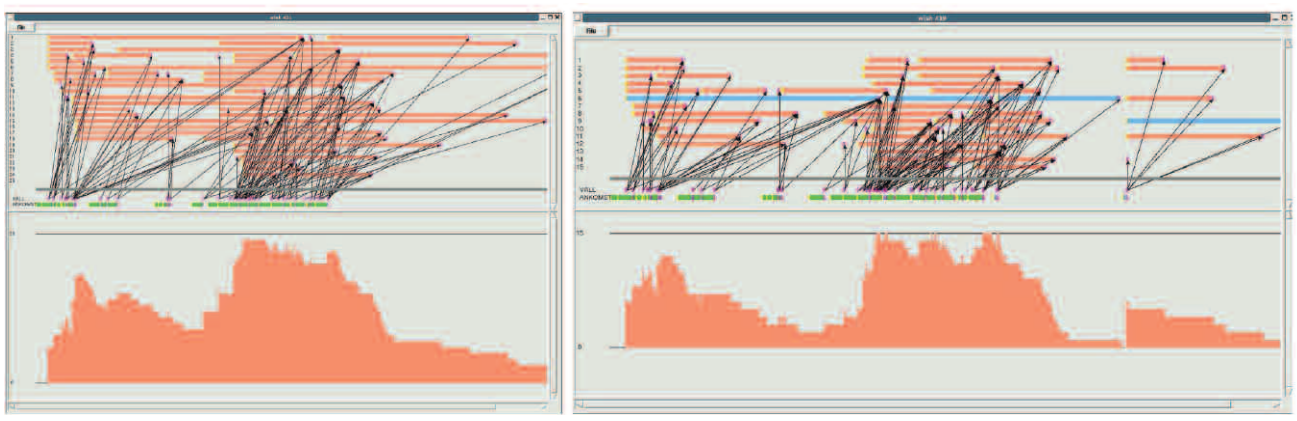

Figure 4: Sävnäs shunting yard without and with temporary trains.

solutions are computed in the order of minutes.

Figure 4 shows the resulting schedules without and with temporary trains. The upper parts are Gantt diagrams of the arrival group (green), barrier (violet) and shunting group tracks (orange). The lower parts show histograms of the number of shunting-group tracks simultaneously being occupied by trains. In the right part of the figure, temporary trains are represented by blue bars. Note how the histogram is spread out (levelled) in time at the cost of a slightly higher load on the entry group and barrier resources used for re-shunting.

Conclusions

The proposed model is quite tentative and despite the fact that so far it has scaled fairly well, it is not clear that constraint programming is the best approach to solve this type of problem. Local search heuristics or an integer programing model may in the end turn out to be more practical. It has also been pointed out to us that the problem has certain similarities with register allocation problems in compiler design.

Link:
http://www.sics.se/iam/

Please contact:
Martin Aronsson and Per Kreuger
SICS, Sweden
Tel: +46 8 6331500
E-mail: Martin.Aronsson@sics.se or
Per.Kreuger@sics.se

The Art of Stacking

by Martin Aronsson and Per Kreuger

Software that optimizes stacking of container ports has the potential to significantly improve transhipment efficiency. We have investigated this problem and present some exciting new approaches to solving it.

Large container ports around the world are major hubs in the global cargo transport system. Efficient management of resources in and around a port is essential since investments in port facilities, vessels and support systems are substantial.

SICS has performed a pilot study to investigate how the efficiency of port operations could be improved. The chosen focus was on the container stacks. SICS developed a general model for the handling of the stacks as stores, a simple demonstration that uses and compares two implementations with different properties.

A container stack is a type of temporary store where containers await further transport by truck, train or vessel. The main efficiency problem for an individual stack is to ensure easy access to containers at the expected time of transfer. Since stacks are 'last–in, last–out', and the cranes used to relocate containers within the stack are heavily used, the stacks must be maintained in a state that minimizes on-demand relocations.

Figure 1: Stack cranes are tight resources in the port.

Stacking Problem Requirements and Objective

Loading and offloading containers on the stack is performed by stack cranes. In order to access a container which is not at the top of its pile, those above it must be relocated. This reduces the productivity of the cranes. Figure 2 shows a scheme of a port indicating the resources involved.

Maximizing the efficiency of this process leads to several requirements. First, each incoming container should be allocated a place in the stack which should be free and supported at the time of arrival. Second, each outgoing container should be easily accessible, and preferably close to its unloading position, at the time of its departure. In addition, the stability of the stack puts certain limits on, for example, differences in heights in adjacent areas, the placement of empty and 'half' containers and so on.

The objective of this work is therefore to plan the movement of the cranes so as to

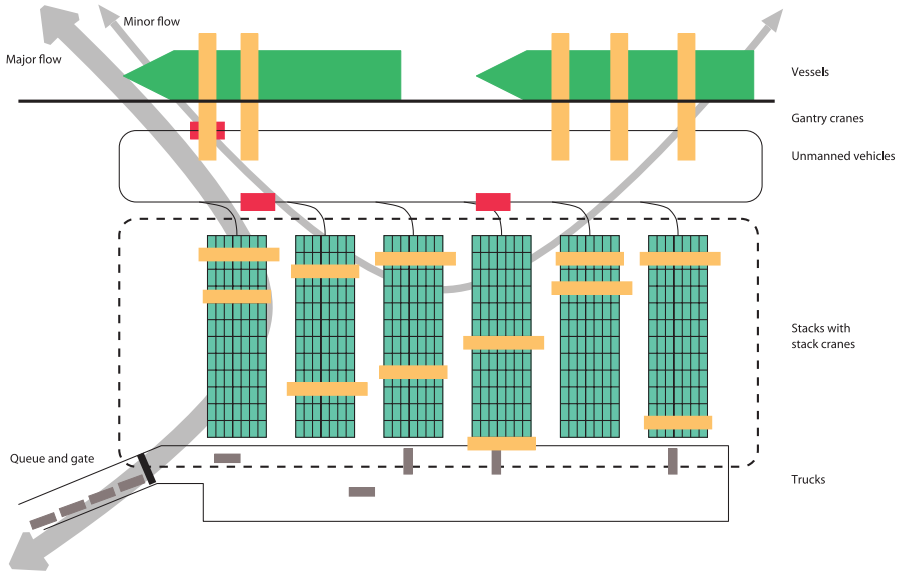

Figure 2: A scheme of a port, where large cranes load and offload cargo from the vessels, unmanned vehicles transport the individual containers between stacks and vessels, and stack cranes are used to load the containers both onto the stacks and from the stacks onto trucks.

fulfil these requirements with a minimum number of movements and/or a minimal waiting time for vehicles and vessels.

Methods for Batch Relocation of Containers

Assuming that each outgoing container in the stack has a fixed departure time, a priority can be assigned to each container. This is essentially the duration in days to its departure.

The following two sections describe two models developed in the pilot study. The first is a constraint programming (CP) model that characterizes the desirable properties of a solution to the relocation problem; the second is a heuristic where the properties of desirable states are instead captured by a more sophisticated cost function than that used in the CP model.

The Constraint Model

This model formulates a number of hard constraints which limit the search for an improved goal state. That is, it formulates the properties of a 'sufficiently' desirable goal state and then searches for a sequence of moves that will take the stack in a direction of such a state. These constraints are:

- *feasibility* - any allocation must be supported, i.e. stand directly on the ground or on top of another container
- *consistency* - only one container can be allocated to each position at any given time
- *priority* - all piles should be sorted so those highest priority containers are always on top of lower or equal priority containers.

When searching for sequences of moves to achieve such a state, only containers with the highest priority and containers above them are considered for moving.

With the addition of a simple cost function, priority can be given to moves that tend to take the containers towards the 'right' area of the stack, clear an area where we expect incoming containers, and/or are quick or 'cheap' to perform.

Local Search

The method is a 'greedy' algorithm that could easily be incorporated into a more general local search mechanism but can also, as in the cases considered in the pilot study, be used directly to improve any given stack configuration.

Let the penalty of a single pile be a weighted sum of its 'unsortedness', the distance of each container in the pile from its ideal x,y-position in the stack and, finally, the distance in its z-position from its ideal z-position. Then, in each iteration:

- choose the pile with the worst contribution to the penalty
- compute the best new placement for the top container of that pile
- if this improves the overall penalty, perform the move and iterate
- if not, choose the second worst pile etc
- when no more improving moves can be found, terminate.

The method is quite sensitive to the exact weights assigned to each factor in the penalty function and, especially for stack configurations with close to full allocation, can easily get stuck in local minima.

Results

Using the ideas described above, the study showed that, even using comparatively simple methods from local search and constraint programming, we could reduce the total penalty for typical stacks by 10-50% depending on, for example, the degree of utilisation of the total stack capacity.

Since the allocation of positions to containers is currently done more or less manually, this has convinced us that it should be possible to achieve significant improvements of lead times, storage utilisation and throughput using improved techniques of the type indicated. Even though the pilot study was based on rather simplistic models, which didn't take into account factors such as stack stability or container content, we are confident that improving the management of the container stacks in a more realistic setting would also yield significant improvements in several of the most important measures of transhipment efficiency.

Link:
http://www.sics.se/iam/

Please contact:
Martin Aronsson and Per Kreuger
SICS, Sweden
Tel: +46 8 6331500
E-mail: Martin.Aronsson@sics.se,
Per.Kreuger@sics.se

Freeway Applications Based on the Macroscopic Model of Traffic Flow

by Tamás Luspay, István Varga and Balázs Kulcsár

The design of modern optimal control strategies is one of the research activities of the Systems and Control Laboratory, SZTAKI, through the project 'Advanced Vehicles and Vehicle-Control Knowledge Centre'. We have applied modern system and control theory to a variety of fields, and this has recently led to new results in urban traffic control and freeway traffic applications

The basis of freeway control is a mathematical description of traffic flow that incorporates as many traffic characteristics as possible. We first summarized and analysed existing models, and found the second-order macroscopic traffic model to be the most accurate and appropriate for our aims. This model is based on the analogy between fluid mechanics and traffic behaviour, and this can be extended to include special traffic phenomena such as congestion. The model works with the traffic variables of density, flow and space mean speed. Although the model equations are precise, there are several unknown parameters that must be properly tuned in order to achieve the best performance.

Model Calibration

We modelled a 4.5 kilometre stretch of a Hungarian freeway with on-ramps and off-ramps and with different geometric properties (ie slope, rising, curves). This allowed us to study the effects of these road properties on the traffic flow. These are reflected in the values of the above mentioned unknown parameters. We studied the determination of the parameters, and worked out three different methods.

The first method is a classical traffic-engineering approach: based on the measurement data, calculate the parameters using the theory of nonlinear kinematic waves. The second is a classical system-engineering approach: based on the measurement data, identify the parameters with nonlinear optimisation theory. Both these methods gave similar results.

During optimisation we also analysed the scope of the model equations. The state limitation of the model was found to be half a kilometre, so our observed section of freeway can be divided into segments. These can be described with equations and then connected to each other according to the boundary relations indicated in the equations.

Freeway traffic estimator

By determining the parameters, we get a compact nonlinear state-space representation of the stretch of freeway. Based on this, algorithms can be developed to improve the capacity of highways. Our main idea was the following: the information provided by traffic flow detectors on freeways is limited to their installation locations. In order to acquire information about conditions between detectors, a freeway traffic estimator is necessary. Since the model equations are highly nonlinear, we chose a well-known nonlinear estimation technique to do this: the Extended Kalman Filter. Our freeway traffic estimator estimates the states – the traffic variables – of each freeway segment, by modelling the noise actuating on the process and disturbing the measurements. The estimator uses real measurement data collected by inductive loop detectors located at the beginning and end of the examined section, and the simulation is evaluated by comparing the estimated variables with real data gathered from a 'control' loop detector located in the middle of the 4,5 km long section. The results show good state estimation under both normal and changing conditions.

The main drawback of these parameter determination methods is that in real traffic, the parameters (and the noise variance) can vary in time and space. Consequently, using constant parameter values is an oversimplification. To overcome this we extended our estimator to give parallel estimation of parameter values and to allow the parameters to vary. This also gives significantly more information about traffic conditions; for example, sharp changes indicate an incident (see figure).

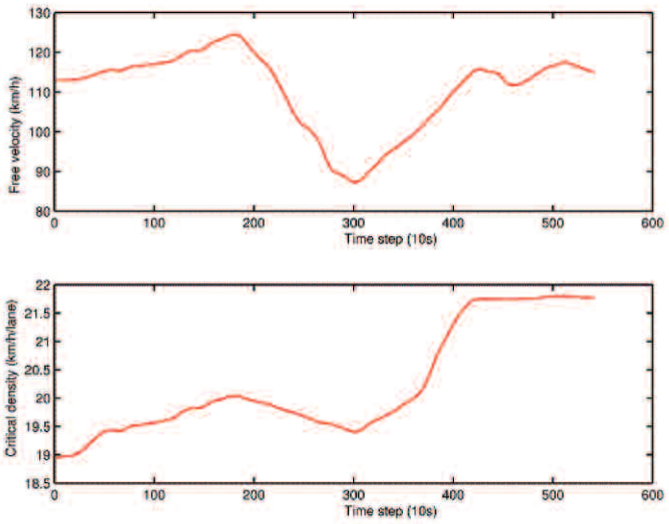

Estimation of important model parameters.

Several applications can now be developed based on this dynamic model and with the usage of the traffic estimator. In our research we designed a new Automatic Incident Detection algorithm to perceive accidents that have occurred on freeways. Our method works with the estimated variables and checks changes in the speed-curve of the section. In this way we can identify accidents between detector stations with an accuracy of 500 metres. The algorithm was successfully tested with

real data. We also drafted the principles of freeway control strategies and simulated them with our model.

Our further research will focus on:
- correction of our dynamic model
- the application of LPV (Linear Parameter Varying) technique to model traffic flow
- MHE (Moving Horizon Estimation) estimation of traffic variables
- adaptive filtering techniques to allow changing noise variance
- developing coordinated freeway control including ramp metering and the usage of variable speed limits
- testing the developed systems in real time.

Traffic control design is one of the main tasks of the long-term project 'Advanced Vehicles and Vehicle-Control Knowledge Centre', managed by SZTAKI, Systems and Control Laboratory in cooperation with the Department of Control and Transport Automation at the Budapest University of Technology and Economics.

Links:
http://www.sztaki.hu/scl
http://www.ejjt.bme.hu

Please contact:
Tamás Luspay and István Varga
SZTAKI, Hungary
Tel: +36 1 279 7266; +36 1 279 6227
E-mail: tluspay@sztaki.hu;
ivarga@sztaki.hu

Performance Evaluation of Heterogeneous Motorized Traffic at Urban Single-Lane Junctions and Roundabouts

by Puspita Deo and Heather J. Ruskin

Managing road networks requires a clear understanding of traffic flow, including causes of congestion or locations of traffic breakdowns. Scientists at Dublin City University are examining the influence of long-vehicle (LV) fraction on urban and inter-urban traffic flow through a road network. A two-component cellular automaton (2-CA) methodology is chosen for heterogeneous motorized traffic. The emphasis is on individual vehicle interactions for a binary mix of vehicle types, namely short (SV) and long vehicles for an urban single-lane junction and roundabout.

Figure 1: Typical traffic in the developing and developed worlds. For example, from left, (a) India;(b) Ireland; (c) a local roundabout.

Traffic in Western or developed countries is characterized by a variety of motorized vehicles such as cars, buses, trucks and motorcycles. This differs from the situation in developing countries, which have diverse mixes of motorized and non-motorized vehicles, including bicycles, cargo tricycles and human- and animal-drawn carts (see Figures 1a and b). Control methods can also vary widely (Figure 1c), and together with the traffic mix, strongly influence urban network flows.

Traffic growth forecasts for Ireland for the period 2002-2040, made by the National Roads Authority (NRA), predict dramatic increases in both LGVs (cars and Light Goods Vehicles combined) and HGVs (Heavy Goods Vehicles). With large traffic volumes, breakdown of traffic flow is likely to occur, particularly on single-lane roads where all vehicles – regardless of whether they are turning or continuing on – wait in a single queue. (For example, see Figure 1c, depicting a local roundabout where few vehicles can be accommodated at any one time.) Nonetheless, most European cities still rely to some extent on single-lane connections to major arterial routes.

Our aim is to address questions relating to the effect of long vehicles on the operation of single-lane junctions and roundabouts. Field observations on traffic patterns can be difficult and time-consuming to obtain, meaning computer simulation models offer a viable alternative for in-depth study and can aid understanding of traffic dynamics. In this context, cellular automata (CA) have emerged as an efficient tool in modelling traffic flow at the individual unit level. Cellular automata approaches are based on discretisation of the continuous variables that describe the traffic flow, in order to define simple vehicle movement rules.

Two-Component Cellular Automata (2-CA) Model for Heterogeneous Motorized Traffic

In this study, heterogeneous motorized traffic flow is modelled at a single-lane junction and roundabout using 2-CA. The state of a road is described in terms of a set of linked cells occupied by at most one particle, corresponding to a standard car of length less than or equal to 7.5 metres.

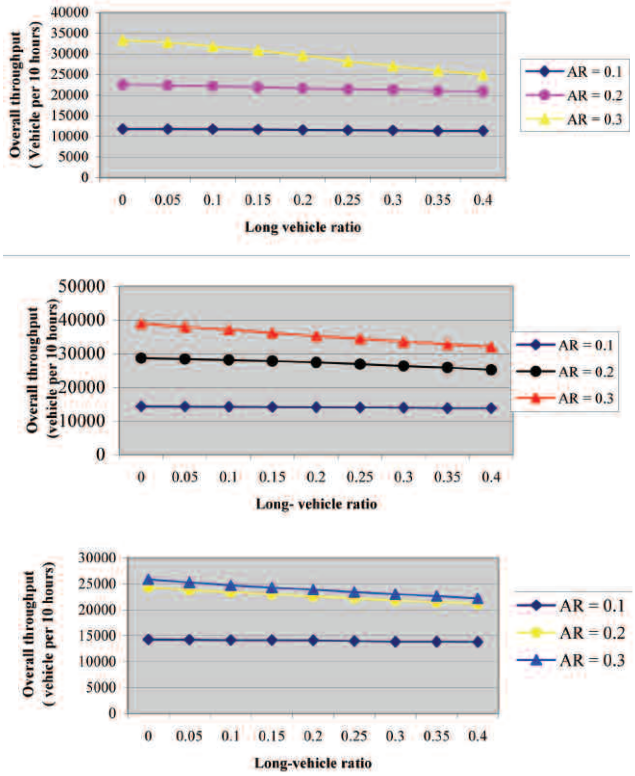

Figure 2: Changes in throughput due to LV:SV mix in through traffic (from top): (a) T-intersection; (b) crossroad intersection; (c) roundabout.

For simplicity, a long vehicle (LV) is taken to be double the length of a standard car (ie two cells are required for one LV). Thus, a short vehicle (SV) is understood to have a length of one cell, while an LV has a length of two. Both SV and LV will move exactly one cell in the next time step if the cell in front is vacant. For intersection and roundabout manoeuvres, an LV requires two time steps and an SV one time step to complete the manoeuvre from entry road onto the intersection or roundabout without causing interruption to flow. The model requires information such as length of entry roads, arrival rate, proportion of each vehicle type, turning rate and simulation time. In this approach, the developed models have been tested with local field data collected manually by Dublin City Council.

Simulations for Road Configurations

The purpose of the study is to ascertain the effect of long vehicles in the traffic mix on the throughput of signalized intersections and roundabouts. Figures 2 and 3 illustrate some results at different types of intersections for varying arrival rates (AR).

Based on these results, a number of points can be made:

- With arrival rates taken to be equal on all entry roads, average overall throughput increases with arrival rate (AR) for homogeneous cases (SV traffic only) up to a critical AR. However, on increasing the proportion of LVs, this critical value is lowered significantly. In all cases (ie T-intersections, crossroad intersections and roundabouts), average throughput is reduced dramatically for large proportions of LVs (see Figure 2).
- An increasing number of LVs in existing traffic streams is expected to worsen the already severe congestion problems present in many cities. From data collected by Dublin City Council, the approximate percentage of LVs is currently around 6%, creating roughly a 12% fall in capacity compared to homogeneous traffic.
- When a network is already operating close to capacity (eg in peak hours), any increase in the proportion of LVs leads to severe bottlenecking and, in extreme cases, gridlock.
- In the simple model described here, additional factors such as longer start-up times are not specifically incorporated, but would also add to delays caused by LV manoeuvres.
- Even simple models can give reasonably good agreement with real data.
- It is clear that investigations into the nature and impact of long vehicles in exclusively single-lane traffic are vital to understanding urban flows. Such shared roads constitute bottlenecks and dictate feeder traffic flow to larger arterial routes; the proportion of LVs in the local traffic has a significant impact on this.
- Only one type of LV is considered here. Clearly, very large buses and trucks will have a different impact on road capacity.

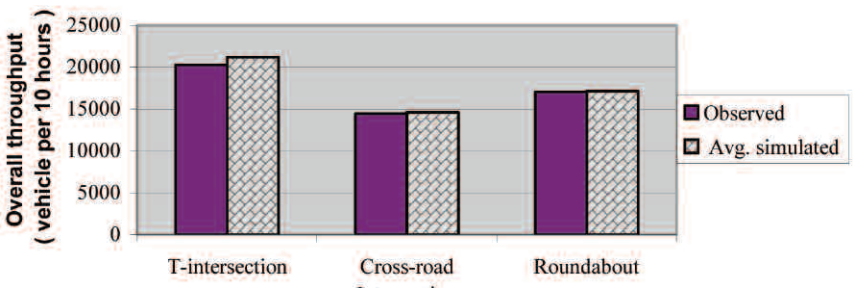

Figure 3: Model validation (comparison of observed and simulated traffic data).

The results of this research suggest that in the case of busy single-lane junctions and roundabouts, restrictions on LV movement will improve throughput to some extent. In the situations studied here, the primary justification for banning LVs is capacity and throughput rather than safety.

Future Work

Studies to date have led to a computationally efficient way of modelling heterogeneous (binary SV:LV mix) motorized traffic flow. Future work will explore vehicular behaviour and other types of LV for a variety of road geometries.

Links:

http://www.springerlink.com/content/274846j686gg0426/

http://www.n3.ie/Transportation/DownloadableDocumentation/file,862,en.PDF

Please contact:
Puspita Deo
Dublin City University, Ireland
E-mail: dpuspita@computing.dcu.ie

Heather J. Ruskin
Dublin City University, Ireland
E-mail: hruskin@computing.dcu.ie

Fluid-Dynamic Approach to Traffic Flow Problems

by Gabriella Bretti, Roberto Natalini and Benedetto Piccoli

A new simulation algorithm based on fluid dynamics, which computes numerical solutions to traffic flow problems on road networks, has been developed. The algorithm reconstructs in real time the evolution of flow and density in road networks, and provides a better approximation with respect to classical models. A satisfactory reconstruction of the network load can be obtained with only a few measurements.

At the Istituto per le Applicazioni del Calcolo (IAC-CNR) in Rome, we have built a simulation tool reproducing the behaviour of traffic within a road network. The simulator prototype can describe the evolution of traffic over time, assigning as input data the initial and boundary conditions on the roads comprising the network, and taking interactions at junctions into consideration. One of the strengths of this tool is the low level of information needed to detect the load on the whole network. This is particularly important for urban traffic control, since actual traffic measurements rely on instruments placed at specific points in the network, typically measuring average fluxes per minute or similar.

There are many motivations for the study of traffic problems, including town planning, investment support, and reduction of congestion, car accidents and pollution. However, modelling traffic is a non-trivial task, and traffic jams in particular are difficult to reproduce.

Vehicular traffic can be treated by different models. We focus on macroscopic fluid-dynamic models, since they mimic phenomena such as the creation of shocks and their backward propagation along roads. Our reference mathematical model is based on conservation laws and represents an extension to road networks of the classical LWR model (Lighthill-Whitham-Richards) for a single road. A road network is a finite set of roads that are modelled by intervals and connected to each other by junctions. It is schematized as an oriented graph, where roads are modelled as arcs and junctions as nodes.

From a numerical point of view, the algorithm is based on approximation methods such as the Godunov scheme and kinetic schemes with suitable boundary conditions at junctions. It has been applied to sample cases and to a variety of real sections of urban networks such as bottlenecks, traffic circles and intersections (see Figures 2 and 3). We have investigated the numerical validity of the approximation algorithm. It has proven to be very fast, providing solutions on networks with a few thousands arcs in less than a second of CPU time on a PC, with ten discretization points per arc.

We then tested the model's effectiveness on the whole road network of the city of Salerno. Using the simulation tool we constructed an algorithm that would trace the position of a car moving into a road network, and give an estimate of the time taken to traverse its route. We also ran tests to reproduce the effects of traffic jams on the trajectory of a single car (Figure 3). In this way, an optimal path can be calculated through a loaded network, taking into account the evolution of traffic over the whole network.

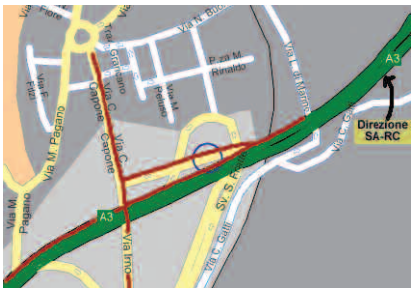

Figure 1: Segment Salerno-Fratte on A3 highway.

Figure 2: A traffic circle – Piazza Re in Rome.

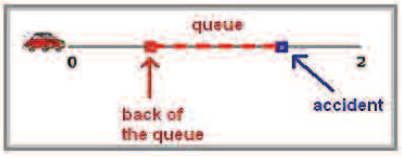

Figure 3: Car moving towards a congestion point.

The fast numerical algorithms described above render the relative optimisation problem treatable. Furthermore, via the acquisition of experimental data measured by sensors located along roads (courtesy of Atac S.p.a. in Rome), we developed a procedure for calibrating traffic parameters on a segment of urban network in the centre of the city of Rome (Figure 4). This procedure is based on the comparison between measured data and the simulated solutions produced by the algorithm, and gives a percentage error of around 19% in the congested period. This is comparable with the average measurement errors of the sensors used and is therefore satisfactory.

The visualisation of simulations is provided by online animation. It is also possible to link the information to a GIS system as shown in Figures 1 and 2.

In the future, we intend to test the calibration procedure on whole networks, and to construct a real-time interface for the automatic representation and visualisation of density curves. Simulations on the network of the city of Salerno were performed in collaboration with the Department of Information Engineering and Applied Mathematics (DIIMA) at the University of Salerno.

Links:
http://www.iac.rm.cnr.it/~bretti/Fluid-dynamic.html
http://www.iac.rm.cnr.it/~natalini
http://www.iac.rm.cnr.it/~piccoli

Please contact:
Gabriella Bretti, IAC-CNR, Italy
Tel: +39 06 88470271
E-mail: g.bretti@iac.cnr.it

Dynamic Routing of Rail Vehicles

by Anders Holst and Markus Bohlin

With the introduction of condition-based and predictive maintenance, today's fixed vehicle cycles designed for fixed maintenance intervals are no longer appropriate. We have shown that it is possible to construct vehicle circuits dynamically in response to condition monitoring and condition counters.

In rail transport, preventive maintenance is typically planned using fixed intervals in units of distance travelled and calendar time. For practical reasons, maintenance is grouped into packages with the same interval. An advantage of fixed interval maintenance is that plans for maintenance can be made in advance with relative precision at the same time as vehicle movements are planned. This is possible because in theory, the vehicle movement plan gives the exact consumption in calendar time and distance travelled for a vehicle at any point in the future.

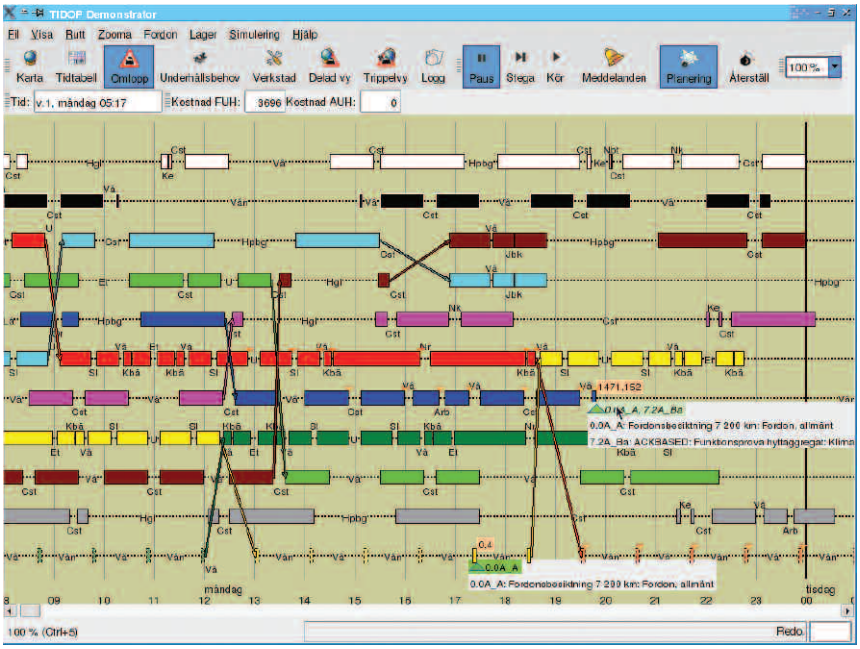

Figure 1: The vehicle circuit view of the prototype software. The operator can either let the computer decide a favourable plan, or can plan vehicular movements manually. The plan is rebuilt automatically to suit the current condition of the vehicles.

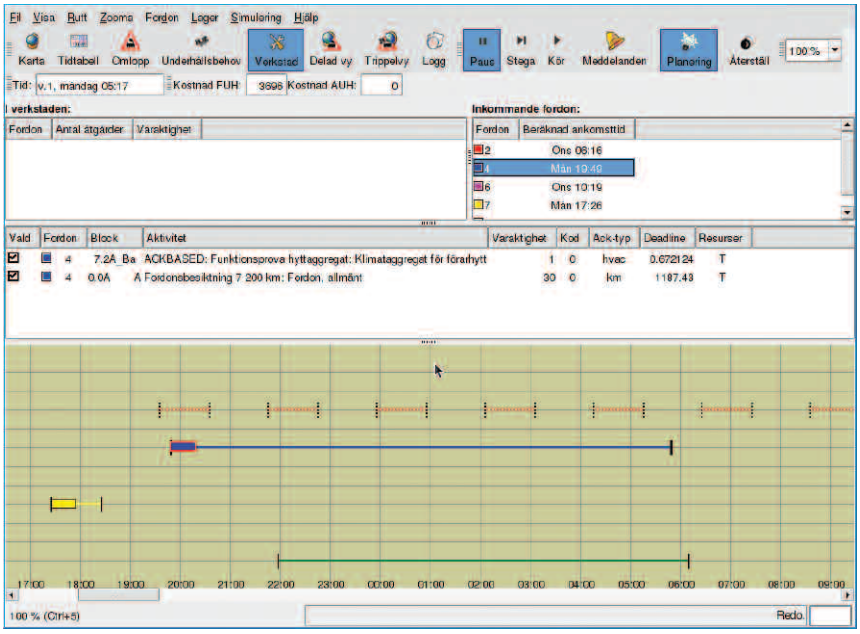

Figure 2: The workshop view of the prototype software. The Gantt chart shows train arrivals and allocated maintenance jobs.

The vehicle movement plan is typically based on fixed vehicle cycles. The cycles are laid out in such a way that each vehicle should be able to pass a workshop at least as often as the most frequent service package requires. Whenever an unplanned service demand occurs, the operator tries to let the vehicle switch routes with another vehicle in order to get it into the workshop as quickly as possible. After maintenance of the faulty vehicle, the two vehicles will typically switch back with each other again, so as not to disturb the vehicle cycles too much.

However, with the introduction of condition-based and predictive maintenance, the use of predetermined service packages and fixed service intervals is no longer appropriate. The drawbacks of such traditional maintenance are that it is inflexible, sub-optimal and – since maintenance is often done unnecessarily and parts are prematurely changed – wastes money. At the same time, other parts may wear out and break before the planned maintenance stop.

To minimize maintenance costs while avoiding expensive breakdowns, trains should be maintained at the best possible moment, considering factors such as vehicle condition, maintenance staff, workshop load, availability of spare parts, co-allocation benefits, downtime costs etc. Because many of these factors constantly change, short-notice planning and optimisation are necessary. Also, since the condition of the vehicles varies considerably with time, the component condition would have to be monitored continuously so that predictions regarding the future status of the vehicle are always available. With this information,

it is possible to recreate with sufficient notice maintenance plans that fit with prediction and at the same time minimize costs.

To illustrate the potential of this concept, prototype software for integrating railroad vehicle diagnosis and operative vehicle circuit planning has been developed at SICS. This prototype combines the advantage of exact knowledge of individual train component conditions with planning algorithms, ensuring that vehicles arrive at the maintenance workshop at the best possible moment. The system simulates and evaluates train movements and wear and tear on components, and dynamically adjusts the vehicle circuits to automatically route vehicles in need of maintenance to the nearest workshop. In addition, spontaneous faults and failures are introduced during the simulation, the appearance of which also triggers re-planning and maintenance stops. The system has proven to be very robust in terms of meeting maintenance deadlines, which naturally is a critical property of a maintenance system.

In short, to take maximum advantage of knowing the exact condition of a technical system, eg a railway vehicle, this information must be connected directly to the company's operative planning strategy. We have shown that such a connection between condition monitoring and operations planning can indeed be established, with very beneficial results.

Link:
http://www.sics.se/iam/index.php?section=timm

Please contact:
Anders Holst and Markus Bohlin
SICS, Sweden
Tel: +46 8 633 1500
E-mail: anders.holst@sics.se, markus.bohlin@sics.se

Statistical Models for Urban Traveller Information and Traffic Management Systems

by Poulicos Prastacos and Yiannis Kamarianakis

Intelligent transportation systems integrate data collection, processing and communications infrastructure with data storage and analytical tools. At IACM-FORTH, researchers in the Regional Analysis Division (RAD) have been working on the statistical modelling issues that emerge in the processing of traffic data streams of great magnitude and wildly varying quality. In particular, they are looking at the detection of traffic sensor malfunctions, imputation of missing or bad data, estimation of velocity and forecasting of travel times in urban arterials.

As vehicular traffic congestion has increased in recent decades, so have efforts to improve data collection, analysis and modelling, and thereby the efficiency of transportation systems. Advanced traveller information systems are now used in cities worldwide. Most of them are based on historical and real-time data, and post current travel-time estimations and future travel-time predictions both on the Web and on message signs at major urban arterials.

Practically every urban traveller information or traffic management system is based on traffic data collected by point sensors (eg single or double loop detectors, overhead video cameras). These provide estimates of flow, volume, occupancy and speed, averaged over time intervals that vary from thirty to ninety seconds. Part of the point-sensor data is usually calibrated with the use of floating car data. Figure 1 depicts a sample from a set of approximately 500 point sensors located in the Athens (Greece) road network. Representative subsets of the database provided by these sensors have been used in testing most of the statistical models developed in RAD.

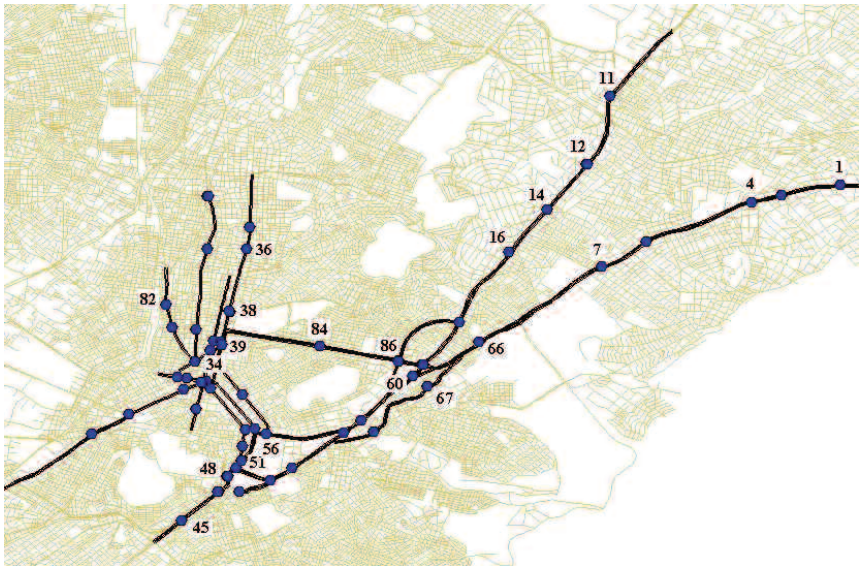

Figure 1: A sample of the point sensors placed in the Athens road network to provide volume, occupancy and speed estimates.

Unfortunately, most of the point sensors employed worldwide are single loop detectors. These can only estimate traffic volume and occupancy and do not directly measure velocity, the most useful variable for traffic control and traveller information systems. Researchers in RAD are using recently developed methods that estimate velocity from single loop data: like a number of related methods these techniques are based on the preceding estimation of mean vehicle length. Moreover, to correct for unstable speed behaviour, espe-

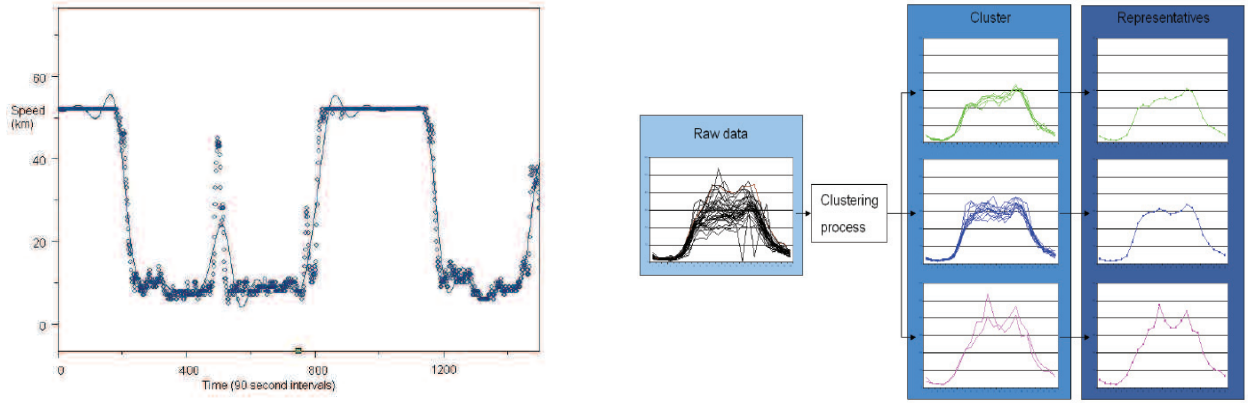

Figure 2: (left) Smoothing spline algorithm applied to speed estimated from a point sensor; (right) curve clustering applied to smoothed point sensor data.

cially during light traffic, smoothing spline algorithms are applied before travel-time estimations (Figure 2, left). To identify locations with similar traffic patterns, curve-clustering methods are applied at an exploratory stage (Figure 2, right). Holes in the data due to missing or bad observations are filled with imputed values; for that purpose, the nonlinear space-time models originally developed for forecasting are employed.

Perhaps the most challenging statistical issue in advanced traveller information systems is travel-time forecasting. Researchers in RAD are pioneers in the development and application of nonlinear space-time models for speed forecasting, and their models display a forecasting performance superior to existing methods. Models are spatial in the sense that they incorporate information related to the road network topology by using suitably defined spatial weight matrices. Models are nonlinear in the sense that relationships between upstream/downstream locations are not constant but are allowed to vary according to time and traffic state (free-flow/congestion). For instance, during free flow an upstream detector may have better predictive power than a downstream one, while the opposite may be true during congestion. Forecasting models are also nonlinear in the sense that error variance is not considered constant: separate autoregressive models provide variable confidence intervals for the forecasts based on time and traffic state.

The current research interests of the RAD group focus mostly on the application of the aforementioned methods to a new traffic information system for Athens. The system will provide real-time information via the Web and cellular phone services. The research is partially supported by LIAISON, a 42-month project co-funded by the European commission under the 6th Framework Programme and led by Alcatel Alenia Space.

Links:
http://www.iacm.forth.gr/regional/
http://liaison.newapplication.it

Please contact:
Poulicos Prastacos
IACM-FORTH, Greece
Tel: +30 2810 391767
E-mail:poulicos@iacm.forth.gr

Mobile Environmental Sensor Systems Across a Grid Environment - the MESSAGE Project

by John Polak

The impact of road traffic on local air quality is of major concern in public policy, and in recent years has stimulated a substantial body of research. This is aimed at improving underlying vehicle and traffic management technology and informing public policy action. Recent work has begun to exploit the capability of a variety of vehicle-based, person-based and infrastructure-based sensor systems to collect real-time data on important aspects of driver and traffic behaviour, vehicle emissions, pollutant dispersion and concentration, and human exposure.

The variety, pervasiveness and scale of these sensor data will increase significantly in the future as sensors become cheaper, smaller and lower in power consumption. This has the potential to improve our understanding of urban air pollution and hence urban air quality. However, handling the vast quantities of real-time data that will be generated by these sensors will be a formidable task. In the case of static sensor grids, significant progress has already been made in the application of e-Science technology to the challenges of real-time data mining and distributed processing.

However, Grids comprising vehicle-based and person-based environmental sensors will be inherently mobile and dynamic in nature, and as such raise significant new challenges. These will require the integration of e-Science with communication and positioning technology, and the linking of all three with domain-specific transport and environmental modelling tools and methods.

Such an integration is the objective of the recently established MESSAGE project (Mobile Environmental Sensor Systems Across a Grid Environment). The project is funded by the UK Engineering and Physical Sciences Research Council and the UK Department for Transport, and involves collaboration between a multidisciplinary team of researchers based at Imperial College London and the Universities of Cambridge, Leeds, Newcastle and Southampton.

There are two main aims of the project. The first is to harness the potential of diverse, low-cost and ubiquitous environmental sensors. These will provide data to address key scientific challenges in the field of transport and environmental monitoring, and modelling and analysis. The second is to develop a flexible and reusable e-Science infrastructure. This will support a wide range of scientific, policy-related and commercial uses and applications for the resultant data, and demonstrate the operation and utility of this infrastructure in a number of case study applications.

The project will develop the capability to equip vehicles and pedestrians with a range of environmental sensors, to track the position of these sensors in the urban environment (using a combination of GPS and wireless network positioning), to retrieve data from them in real time, to integrate these data with those from more conventional (static) environmental and traffic sensors and to estimate pollutant concentrations and individual exposure. The project will build on and extend a number of outputs from the UK e-Science programme, including the DiscoveryNet project (http://www.discovery-on-the.net) and the OGSA-DAI and OGSA-DQP projects (http://www.ogsadai.org.uk), as well as taking advantage of extensive experience across the consortium in measuring and modelling the environmental impacts of traffic.

The case-study applications will be carried out in London, Cambridge, Gateshead and Leicester. They will build on the Universities' existing collaborative arrangements with the relevant local authorities in each site, drawing on substantial existing data resources, sensor networks and ongoing EPSRC and industrially funded research activities. These applications will address important problems that to date have been difficult or impossible to solve due to a lack of relevant data. These problems are of three main types: (i) measuring human exposure to pollutants, (ii) the validation of various detailed models of traffic behaviour and pollutant emission and dispersion, and (iii) the development of transport network management and control strategies that take account not just of traffic but also air quality impacts. The various case studies will look at different aspects of these questions and use a variety of different types of sensor system to do so. In particular, existing sensor networks in each city will be enhanced by the selective deployment of a number of new sensor types (both roadside and on-vehicle/person) to increase the diversity of sensor inputs.

Links:
http://www.message-project.org
http://www.imperial.ac.uk/cts

Please contact:
John Polak
Imperial College London, UK
E-mail:j.polak@imperial.ac.uk

Statistical Physics Algorithms for Traffic Reconstruction

by Arnaud de La Fortelle, Jean-Marc Lasgouttes and Cyril Furtlehner

Concepts and techniques from statistical physics have inspired a new method for traffic prediction. This method is particularly suitable in settings where the only information available is floating car data. We propose a system, based on the Ising model from statistical physics, which both reconstructs and predicts the traffic in real time using a message-passing algorithm.

The idea behind this development is that current models, while suited to traffic reconstruction and prediction on a motorway, have severe drawbacks in other areas. In the first place, the information available is heterogeneous: sensors can be magnetic loops, video cameras or floating car data (the data retrieved and sent by a car to a server). This generates a lot of noise for analysis.

Second, while some existing urban and inter-urban areas have traffic management and advice systems that collect and analyse data from stationary sensors, many do not. This is the case particularly in rural areas, where crashes account for more than 60% of all road fatalities recorded in OECD countries. Hence, the need for a system that can cover these roads is compelling if a significant reduction in traffic-related deaths is to be achieved.

The Stochastic Model

Most current traffic models are deterministic, that is, described at a macroscopic level by a set of differential equations linking variables such as flow and density. Such models are quite well adapted and efficient on motorways where a fluid approximation of the traffic is reasonable. However, they tend to fail for cities or rural roads, since the velocity flow field is subject to much greater fluctuations. These are induced by the nature of the network (presence of intersections and short distance between two intersections) rather than by the traffic itself. This requires a stochastic model.

These considerations have led to a hybrid approach, taking full advantage of the statistical nature of the information. In order to reconstruct the traffic and make predictions, we propose a model – the Bethe approximation – to encode the statistical fluctuations and stochastic evolution of the traffic, and an algorithm – the belief propagation algo-

rithm – to decode the information. These concepts are familiar to the computer science and statistical physics communities, since it has been shown that the output of belief propagation is in general the Bethe approximation.

The model is shown in the figure. The network of roads is classically represented as a graph, and traffic state on an edge (a road) is described by a binary value: 0 for fluid and 1 for congested. The important feature of the model is that traffic states are correlated only through the coupling of neighbours. This is the Bethe approximation, a method for finding the exact solutions of certain quantum many-body models. A model such as the Ising model, a simple model used in statistical mechanics, displays a phase transition phenomenon with respect to the value of the coupling. From the point of view of a traffic network, this means that the model is able to describe binary traffic regimes on the whole network: either fluid (most of the spins up) or congested (most of the spins down). This represents the traffic quite accurately: when one part of a road network is congested, it is common that all parts are also congested (and vice versa).

The algorithm used for reconstruction, the belief propagation, reconstructs a Bethe approximation from real data. In fact, the data collected from the probe vehicles is used in two different ways. First, data is collected over long periods in order to estimate the model, ie matching the correlations with historical data. This operation is expensive but can be done once, and updated only if the general behaviour of the network changes. Second, at every period of prediction refreshing (typically five minutes), a reconstruction is performed to match the current data with the model, using the correlations calculated first. This leads both to reconstruction and prediction, since we use a space-time graph.

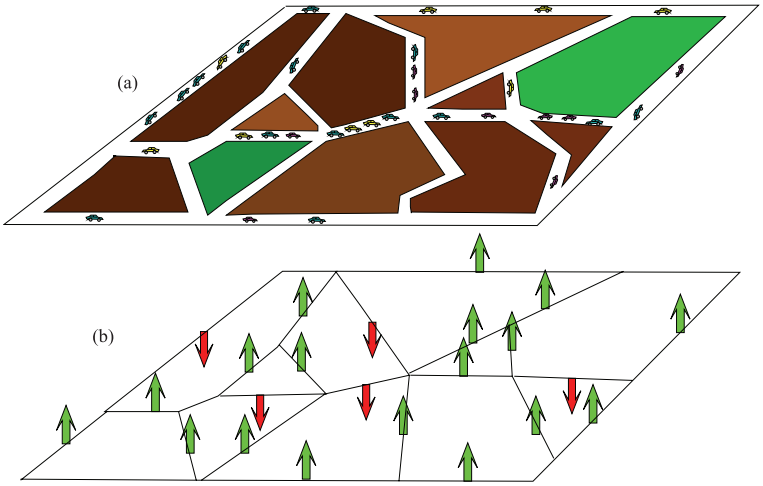

Traffic is modelled as a binary-valued graph.

This algorithm has been implemented, and initial tests using a traffic simulator to generate traffic data show that it is fast (real time for a medium network). The main issue now is stability and precision. The next step is therefore to test it with real data.

Please contact:
Arnaud de La Fortelle
Ecole des Mines de Paris, France
E-mail: arnaud.de_la_fortelle@ensmp.fr

Jean-Marc Lasgouttes
and Cyril Furtlehner
INRIA, France
E-mail: {Jean-Marc.Lasgouttes,
Cyril.Furtlehner}@inria.fr

From Traffic Prediction to Collaborative Navigation Solutions

by Arnaud de La Fortelle, Angel Talamona and Mikaël Kais

The research collaboration LaRA (La Route Automatisé) between the IMARA (Informatics, Mathematics and Automation for La Route Automatisée) Lab at INRIA and the 'Centre de CAO et Robotique' CAOR at the Ecoles de Mines de Paris, improves the state of the art of traffic prediction technology through statistical modelling based on instrumented probe vehicles.

Working together with European industrial and research partners like TNO, Motorola and Intempora, the European project REACT (Realizing Enhanced Safety and Efficiency in European Road Transport) has demonstrated breakthrough technology based on integrated in-car sensing and a traffic management centre. SENDA, an INRIA spin-off company that is licensing LaRA's traffic modelling technology, has built on the work of REACT to overcome technical issues for the cooperative system's deployment.

The goal of cooperative systems like REACT is to increase the efficiency of European road transport by providing drivers with up-to-date information on road conditions. For this, vehicles sense real-time natural and infrastructure conditions both within themselves and in their vicinity, and transmit these data to a central server. There they are aggregated and analysed by an integrated set of models to generate both vehicle-specific recommendations and information for relevant authorities.

Analysing Traffic with Probe Vehicles to Generate Customized Road Advice

The analysis of the traffic conditions and safety risks of different routes can only be performed by a central server,

which collects and analyses data coming from a large number of geographically distributed vehicles. REACT has key advantages over current systems, including:

- mobile rather than stationary sensors, covering all routes where vehicles travel
- measurement of relevant natural and infrastructure conditions
- customized, model-based recommendations transmitted to individual vehicles.

The REACT technology has been integrated into a fleet of instrumented cars, used by LaRA for cooperative vehicle technology research (see Figure 1).

Figure 1: REACT sensor vehicles (4 LaRA C3).

Technology Transfer Issues

While a proof of concept might be enough to validate research activities, actual innovation, in the sense of the economic exploitation of new technology, needs to go further. There are three main lessons to be learnt from the REACT project. First, a large number of players is necessary to put in place a running cooperative system based on probe vehicles (eg public road operators, software companies, equipment providers), and the aims and motives of these parties may not be the same. Second, the integration and successful communication of a 'proof of concept' initiative is an excellent tool for making these concepts comprehensible and encouraging stakeholders to reach operational solutions. Third, a relatively high threshold of deployment is necessary (3-5 % of vehicles in the entire road network must be equipped) to make traffic technology based on probe vehicles exploitable.

SENDA's Alternative, Internet Personal Navigation Device for Collaborative Navigation

SENDA is a mobility-oriented software company and a spin-off from INRIA. It licenses LaRA's statistical model for traffic prediction based on probe vehicles, which was developed during the research project REACT. Combining the technical constraints for the deployment of such breakthrough technology with the need to reduce the number of players in the value chain, SENDA conceived an alternative approach: a collaborative GPS navigation device. This high-end navigation device corresponds to European drivers' demands for improved route guidance, based on better-quality road traffic information and higher coverage of the road network. The principle is simple: the GPS navigator is supplied with an embedded wireless Internet interface and the ability to act as a mobile traffic sensor. The navigator's position and speed are anonymously transmitted to a central management server that statistically processes this information and then consolidates it with the data coming from the road operator's traffic information centre

This concept for innovative navigation is derived from INRIA's previous work in the European projects Cybercars and Ozone, and from the SENDA prescriptive navigation. The latter was a finalist in the European Commission's GST Service Submission Contest. SENDA's first collaborative navigator prototype was introduced during the ITS World Conference 2006. Commercial vehicle renting companies are already interested in such an approach and are discussing with SENDA the implementation of high-end GPS navigation pilot devices for professional drivers.

SENDA has exploited LaRA's research to produce innovations in traffic prediction technology for the next generation of traffic management systems. The REACT experiments helped SENDA move ahead on their reflections, and synthesize a convenient product targeting an emerging market. Deployment of traffic prediction solutions by SENDA will provide valuable real-world feedback to LaRA's research, which in turn will lead to improvements in and testing of predictive models.

Link:
http://www.react-project.org/

Please contact:
Arnaud de La Fortelle
Ecole des Mines de Paris, France
E-mail:
arnaud.de_la_fortelle@ensmp.fr

Angel Talamona, Mikaël Kais
SENDA, France
E-mail:
Angel.Talamona@senda-online.com,
Mikael.Kais@senda-online.com

Figure 2: SENDA's personal navigation device for collaborative navigation.

Logistics Planning Using the POEM Language System

by Jianyang Zhou, Jianyi Zhou, Jiande Zhou, Yann Courtet, Weiping Wu, Hua Ni

In modern logistics, multi-modal transport planning and vehicle routing represent two major problems. Within a supply chain, the first determines how to forward freight between different regions, while the second plans the local fleet of vehicles, both in order to minimize transport costs. These problems are significant for third- and/or fourth-party logistics. In practice, complex transport conditions such as on-time delivery, capacity (quantity/volume/weight), work regulation (time/break/makespan), transport mode, security and coupling constraints may lead to computational challenges for the optimisation engine.

POEM (Programming in Operational and Expressive Models) is a planning and scheduling solution-building system. The POEM Solution Builder is composed of a core technology and an auxiliary one: Natural Constraint Language (NCL) and Visual Script.

NCL is a descriptive language for business modelling and problem-solving based on conventional mathematical logic. Visual Script in POEM is a GUI script that supports solution visualisation (basic graphics/map/time) and interaction management. In an all-in-one style, POEM embeds OR algorithms, logic, heuristic rules and GUI in a single system to offer users a cost-efficient choice with benefits in training, management, development, deployment and maintenance.

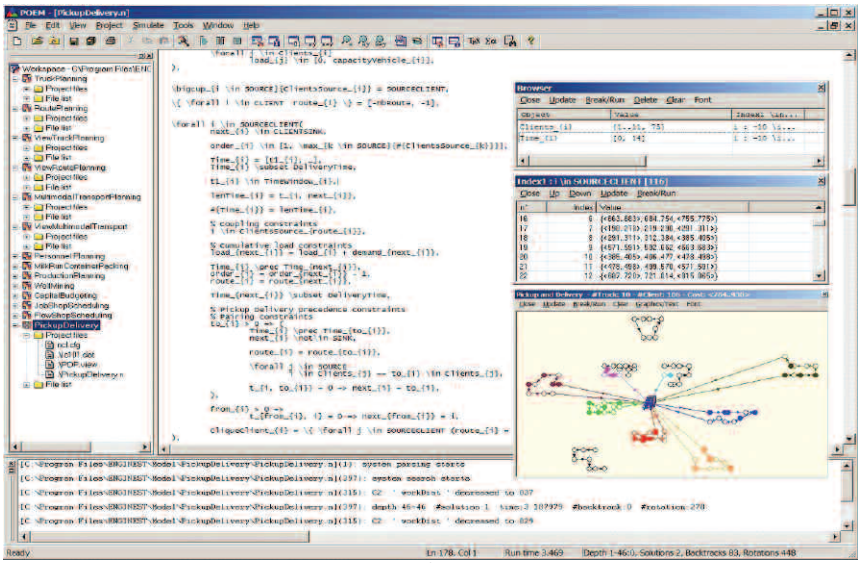

Figure 1: POEM Solution Builder - modelling and diagnosing in a visual manner.

As a development tool, POEM narrows the gap between developers and end users thanks to its visual modelling and debugging interface. Figure 1 shows the visual nature of POEM modelling and diagnosing.

POEM is also designed to be an application platform. In order that a single solver can support a variety of applications, POEM advocates the development of modular solutions for different businesses. For instance, production planning and scheduling, distribution optimisation and personnel planning can all be performed on the POEM platform. To illustrate this point, two POEM solutions for logistics planning are presented below.

Multi-Modal Transport Planning

Multi-Modal Transport Planning (MMTP) involves planning the flow of goods and services along a supply chain. For example, fourth-party logistics requires that transport costs be optimized for on-time delivery of goods from vendors to customers. Transport occurs through multi-modal trade lanes, such as air, rail, road and shipping, with each of these incurring a different cost. In addition, a feasible transport plan must simultaneously satisfy constraints such as vendor (production) capacity, vehicle (route) capacity, transportation schedules (eg an air company may offer several flights per day), and cross docking and service times. Everyday examples of MMTP applications include mail/parcel delivery, freight forwarding and so forth.

Compared to vehicle routing, MMTP is a capacitated freight-routing problem. Each route may have several shifts per day; for example, air transportation may have one flight at 9:00am and a second at 2:00pm each day, with each flight having a limited capacity (quantity/volume/weight). For each customer order, freight routing determines an itinerary from origin to destination that identifies each form of transport used and the overall schedule, and satisfies all the delivery constraints. See Figure 2 for an illustration of multi-modal transport plans.

Vehicle Routing and Scheduling

Vehicle routing and scheduling is the optimisation of delivery routes, which may be subject to constraints as diverse as vehicle capacity (quantity/volume/weight), work regulations (time/break/makespan), time windows (eg a customer may prefer a morning to an afternoon delivery), distance, route accessibility and coupling (ie certain customers should or should not be serviced by the same truck). There may also be multiple objectives: for example, to minimize the number of vehicles, worktime, makespan and mileage.

In real life, vehicle loading can be a significant concern for logistics companies: trucks are expected to be fully loaded for most of their routes (see the

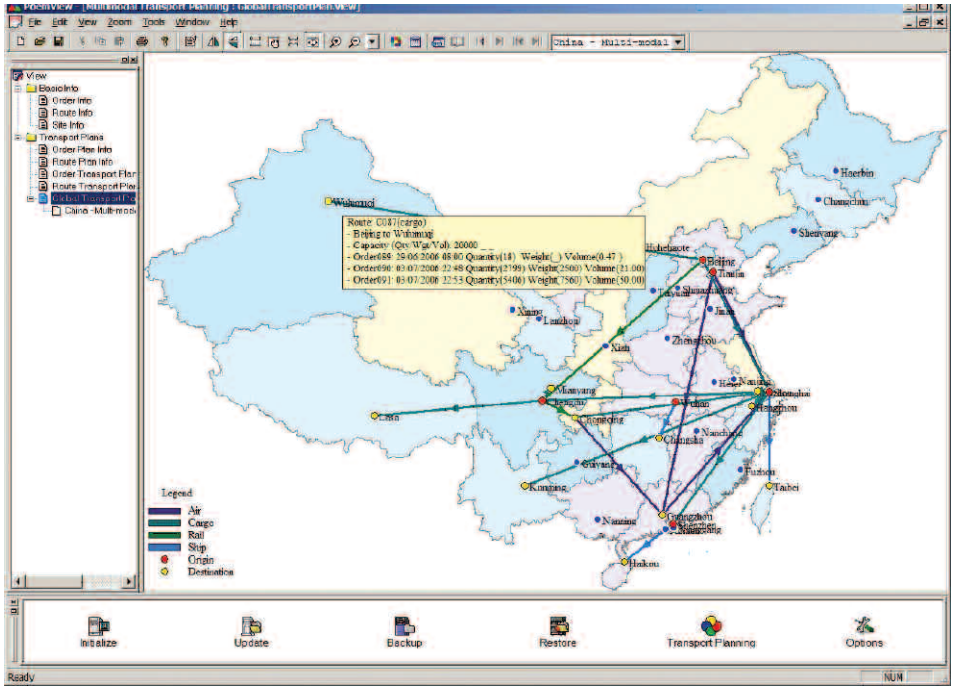

Figure 2: Multi-Modal Transport Planning (example for 200 customer orders).

statistics in Figure 3: almost all the trucks are well loaded). This makes it necessary to first work out the truck plans so as to balance truck loading and take into account time constraints. After truck planning, the system computes the best route and schedule for each trip subject to work regulations and delivery constraints.

Re-Planning and Re-Scheduling

For practical solutions, planning flexibility is of great importance. Circumstances will arise where it is necessary to modify transport modes, delivery schedules, or to shift customer orders from one route to another. In other words, operators may need to intervene in the system's planning solution based on their expertise or experience. For both freight and vehicle routing, the system allows flexible re-planning and re-scheduling to meet the requirements of customer changes and interaction.

This research has been undertaken by ENGINEST. The company is a spin-off from INRIA and is working in the area of logic programming and operations research, with particular interest in resource optimisation for logistics, manufacturing, and human resources.

The 'POEM Modular Solutions' project was started in 2000, and is being extended to better deal with industrial applications in different domains. Future work will look into more powerful OR algorithms and more user-friendly interfaces for the modular solutions. Cooperation with renowned laboratories in these areas is expected.

Link:
http://www.enginest.com

Please contact:
Jianyang Zhou
Enginest, France
Tel: +33 3 83 55 12 99
E-mail: zhou@enginest.com

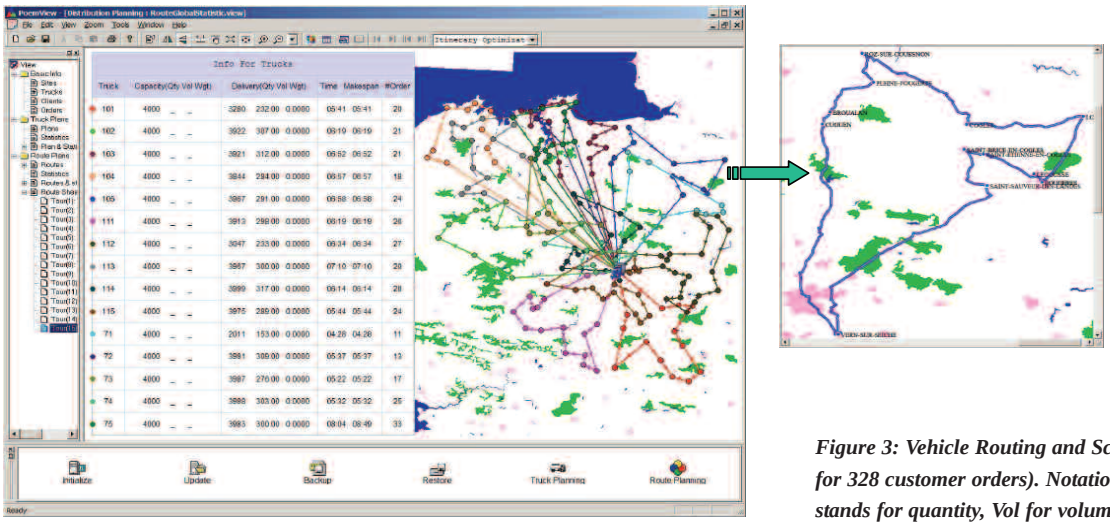

Figure 3: Vehicle Routing and Scheduling (example for 328 customer orders). Notation explanation: Qty stands for quantity, Vol for volume, Wgt for weight

Using GIS for Optimisation in Transportation Planning

by Günter Kiechle

The use of Geographic Information Systems (GIS) for preparing input data for optimisation algorithms improves the practical applicability of such algorithms in the field of transportation planning. This promising combination of technologies is the subject of a collaboration between the University of Vienna and Salzburg Research.

Vehicle Routing Problems (VRPs) are a widely investigated class of problems in combinatorial optimisation, and include many transportation tasks (eg parcel services). In general, a VRP consists of a set of customers that must be served via a fleet of vehicles, each of which leaves from and returns to a central depot. The type of VRP determines whether customers have goods delivered to them, are transported from one location to another, or are served in some other way.

Using GIS for Real-World Input Data

In research, most solution techniques for this class of problem are designed and tested by means of synthetic problem structures. However, the tackling of real-world VRPs requires a thoroughly elaborated data basis in order to provide reasonable outcomes. If this is not the case, even the best solution techniques are of no use for practical applications.

Essential input data for real-world VRPs is gathered by using Geographic Information Systems (GIS). Whereas most researchers use Euclidean distances between customers and depots for their optimisation algorithms, a GIS can provide real distance information derived from a digital road network.

Experiences in former projects showed that using distance data of limited quality in optimisation algorithms leads to results of even more limited quality. In the worst case, a valid solution for a given input dataset might actually be unfeasible in reality.

To get distance information of sufficient quality, the most detailed street network commercially available for the considered region should be used. Unfortunately static distance information from a digital road network does not correlate directly with real travel times because of dynamic influences like traffic jams, road works and weather conditions. Travel times also depend on parameters such as driving style and vehicle type, which are particularly hard to quantify.

Empirical Evaluation of Travel Times

In a former project, the quality of calculated driving times was evaluated by comparing them with actual driving times for 91 trips within a range of 20 to 250 kilometres in eastern Austria. The average, maximum, minimum and standard deviation of actual minus calculacted driving times in minutes recorded by tachographs is given in Table 1.

Deviation	minutes
Average	0.6
Minimum	-32
Maximum	24
Standard	10.3

Table 1: Deviation of driving times in minutes.

This analysis comprised trips to more than forty locations carried out by more than twenty drivers on several days and at all times of day. All the vehicles were of the same type and were able to exceed the allowed speed limit on all the roads used. These results enable us to specify safety margins that are suitable to counteract these observed variations in driving time. The results of this evaluation were used to define reasonable buffer values for optimisation procedures in order to achieve highly robust solutions for use in practical scenarios.

Integration of GIS and Optimisation

In general, optimisation algorithms are implemented in a highly separated software component and used via a well-defined interface. GIS are not only needed for input data processing but in some cases also for enclosing the user interface of a decision support system that guides decision makers by allowing them to schedule vehicles visually. In the following example, GIS plays an integrating role by handling user interactions as a graphical user interface on the one hand and managing communication between all other required software components on the other. A typical example of system architecture is depicted in Figure 1.

The system works as follows: after retrieving customer coordinates and other relevant data from the CRM-Application, the Front-End uses the Network Toolbox to obtain a distance matrix containing travel times for all possible customer pairs. The Front-End forwards all information to the Optimisation Module, which applies the optimisation algorithm. Resulting tours are returned to the Front-End and given a geographical representation using the Network Toolbox. Finally, a complete list of travel plans including customers, street paths and visiting order for the requested period of time may be generated and presented to the application user.

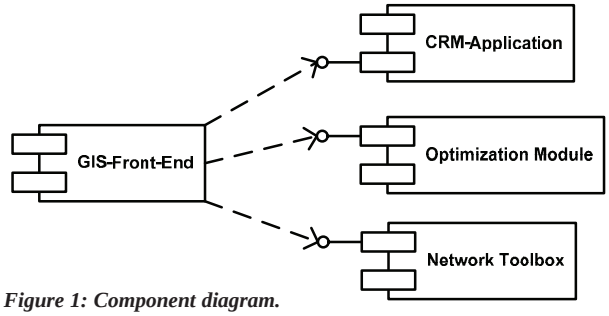

Figure 1: Component diagram.

Implementing the Optimisation Algorithm

Vehicle Routing Problems belong to the group of NP-complete programs that are known to be hard to solve. In general, the computation time for solving a VRP increases exponentially with the overall number of customers. Even execution times for heuristics and meta-heuristics suffer from this effect. For optimisation problems of a certain size only the development and implementation of highly parallel algorithms may achieve reasonable execution times.

Optimisation algorithms are usually implemented in Fortran or C due to the high-performance compilers that are available for these languages. Next to efficiency, additional criteria such as programming convenience, adoption of technology and availability of tools could also justify the use of C++ as a programming language.

Real World VRP Example

The use of GIS for optimisation in transportation planning described above was applied in a project dealing with the development of a decision support system for transportation of blood donations in Austria. The underlying optimisation problem originates from the blood collection process of the Austrian Red Ccross blood program, where processing requirements state that all blood must be processed in one centrally located blood bank within four hours of donation. This restriction causes a significant amount of vehicle movement, which could be reduced with problem-tailored optimisation algorithms by more than 25%, as measured in total driving time. Figure 2 depicts the user interface of the transportation-planning tool developed in this project.

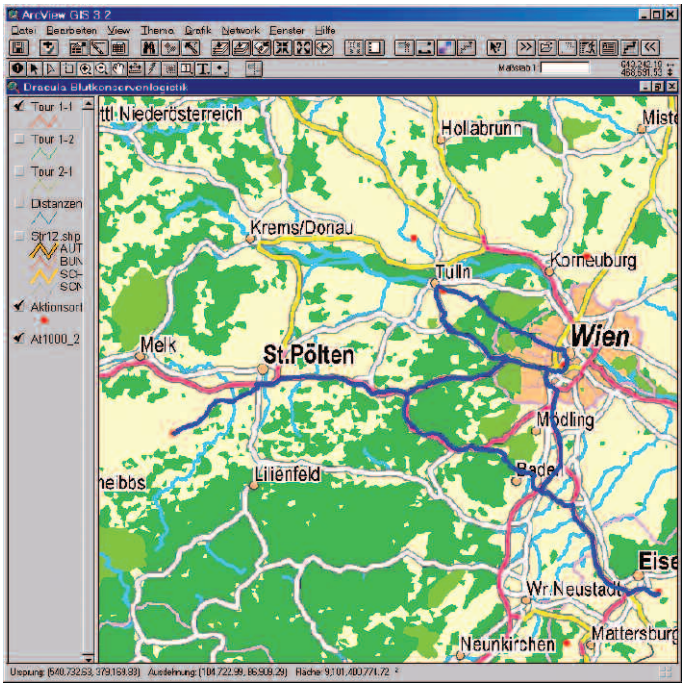

Figure 2: User interface of a decision support system for transportation planning.

Please contact:
Günter Kiechle
Salzburg Research
Forschungsgesellschaft, Austria
Tel: +43 662 2288 421
E-mail:
guenter.kiechle@salzburgresearch.at

GIS Technology for Maritime Traffic Systems

by Cyril Ray, Thomas Devogele, Valérie Noyon, Mathieu Petit, Sebastien Fournier and Christophe Claramunt

At the Naval Academy Research Institute in France, collaborative research in the fields of Geographical Information Systems (GIS) and maritime transportation is producing data models and computing architectures that favour the development of traffic monitoring and analysis for decision-aid systems.

Safety and security are constant concerns in maritime navigation, not least because of the constant increase in maritime traffic and reduction of crews on decks. This has favoured the development of automated monitoring systems such as the Automatic Identification System (AIS) and the Electronic Chart Display and Information System (ECDIS). However, officers on the watch and monitoring authorities require additional and advanced decision-aid solutions that will take advantage of these communication and cartographical systems.

The development of geolocalisation, information and telecommunication technology offers higher traffic data availability, but presents new challenges for the integration, analysis and delivery of maritime traffic data. Hence there is a great need for integrated traffic systems that are able to cope with all these technological domains. The development of a real-time integrated maritime platform implies a reconsideration of storage, modelling, manipulation, analysis and visualisation functions, since current models have not been designed to handle such dynamic phenomena.

One of the most promising solutions to this problem is the integration of Geographical Information Systems (GIS) with maritime navigation systems. Our approach relies on a monitoring project, the purpose of which is to integrate heterogeneous positional data from three sources: an Automatic Identification System (AIS), public traffic data available from the Internet, and a real-time monitoring system developed for sailing races. These systems generate geolocalisation frames from different types (eg AIS frames) and from different ad-hoc networks (eg VHF for AIS, WiMAX for

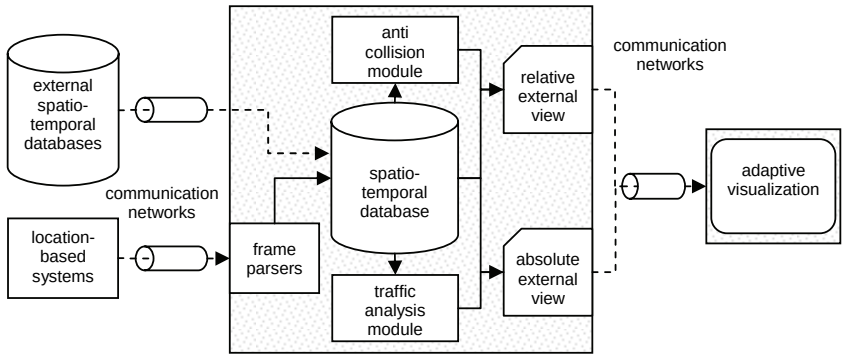

Figure 1: Integrated GIS.

sailing races). Coupled with external databases, this constitutes the input for an internal spatio-temporal database that manages historic and real-time data. These data include maritime trajectories, vessel characteristics and environmental data.

The framework developed so far integrates several modules. One of these is an anti-collision function that monitors the risk of running aground and evasive ship behaviour. This module also integrates simulation capabilities in order to control and predict the evolution of ship behaviour and trajectories. These simulations are based on a multi-agent system and micro-simulation capabilities, where ships are modelled as autonomous agents acting in their environment according to maritime rules.
The module is designed for maritime authorities and for educational and training purposes. Moreover, the traffic analysis module integrates intelligent inference mechanisms that can use data mining to derive traffic patterns. The objective is to observe and understand maritime traffic at different levels of granularity.

Two complementary external views are of interest in maritime traffic monitoring. One is the conventional absolute view of spatial trajectories. The other combines the relative position and velocity of mobile ships with respect to an observer. This relative external view helps users to perceive traffic evolution according to their point of view and aids their perception of the underlying processes that emerge from the ships' behaviour.

Current research into users' interaction with the system is focusing on the concept of adaptive GIS. This can be defined as a generic and context-aware GIS that automatically adapts according to its context. The context is defined by (i) the properties and location of the geographical data being manipulated (eg maritime traffic data); (ii) the underlying categories that reflect different user profiles (eg port authorities); and (iii) the characteristics of the computing system, including Web and wireless techniques. In all, the use of adaptive GIS as a decision-aid system appears to be a useful approach for maritime transportation systems.

The framework developed so far represents a preliminary step towards the development of integrated GIS and maritime information systems. A number of research challenges remain. One is the development of cross-domain protocols and exchange standards for the transmission and interoperability of traffic data; another is the integration of different geographical information sources that are combined, adapted and shared in real time between different levels of users acting in the maritime environment. Conventional statistical, geographical data analysis and visualisation methods should also be adapted to the specific nature of maritime traffic information. Traffic data and applications are usually physically allocated in different geographical locations and are computationally expensive in terms of the data volumes generated. Hence, there is also a need at the implementation level for the development of a GIS-based distributed computing environment, and computational and processing capabilities. Overall, the development of GIS for maritime information systems should improve the management and planning of maritime navigation, thus favouring the search for a safer sea.

Links:
GIS research group of Naval academy:
http://www.ecole-navale.fr/fr/irenav/groupes/sig/sig-en/

Please contact:
Cyril Ray
Naval Academy Research Institute, France
Tel: +33 298233611
E-mail: ray@ecole-navale.fr

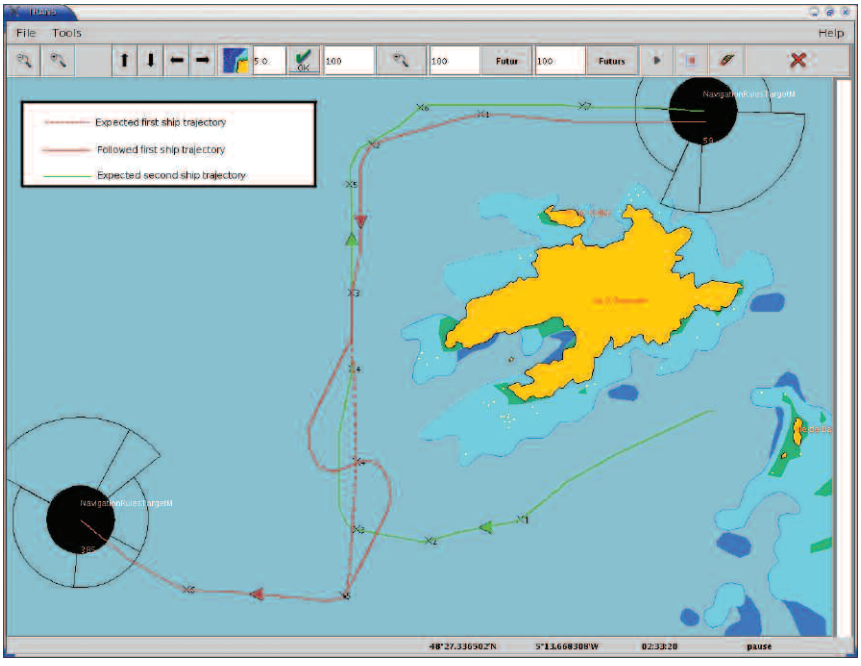

Figure 2: Illustration of anti-collision module.

PeerMart: Decentralized Auctions for Bandwidth Trading on Demand

by David Hausheer and Burkhard Stiller

PeerMart defines a fully decentralized auction-based marketplace layered on top of a peer-to-peer (P2P) network, and makes the trading of services over the Internet technically and economically feasible. While being generally applicable to any service-trading scenario, PeerMart shows great potential as a new method for scalable and reliable bandwidth trading on demand.

The Internet is becoming increasingly popular as an electronic marketplace. As a result of rapid technological progress, the near future will show a wide variety of goods and services being traded over the Internet by huge numbers of buyers and sellers. Such a vast marketplace requires reliable trading mechanisms that are truly efficient and scalable.

At the same time, peer-to-peer (P2P) networks are emerging as a new design approach for building scalable and fault-tolerant applications. By means of resource aggregation and replication, P2P-based applications benefit from peer resources that otherwise lie unused, and thus provide much better performance and greater robustness than traditional client/server-based applications.

PeerMart Architecture and Design

PeerMart is a fully decentralized auction-based marketplace layered on top of a P2P network (see Figure 1). It combines the economic efficiency of double auctions with the technical performance and scalability of structured P2P networks. The solution supersedes the need for a central auctioneer by distributing broker functionality over all peers. Redundancy is used to achieve robustness even in the presence of malicious peers.

The auction-based pricing mechanism is complemented by a decentralized accounting scheme called PeerMint, which provides accountability for applications in a secure and scalable manner. The scheme uses multiple distributed account holders and session mediation peers to store and update accounting information on individual peers and sessions. It also enables peers to be charged for service usage in an aggregated manner.

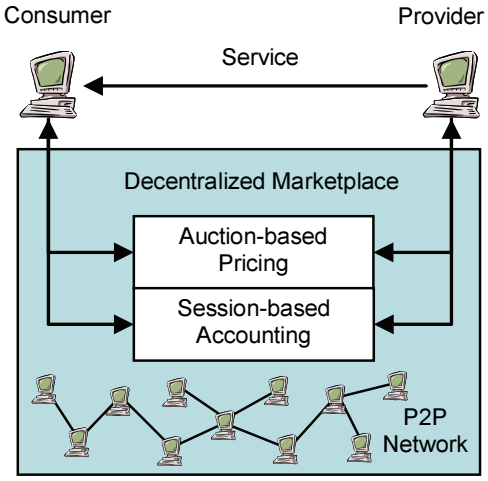

Figure 1: PeerMart – a fully decentralized P2P-based marketplace.

David Hausheer and his advisor Prof. Burkhard Stiller at the Institute TIK at ETH Zurich developed PeerMart's fully decentralized market approach during Hausheer's PhD candidacy. The work was completed in 2005. A prototype of the design has been implemented in Java using the FreePastry structured P2P overlay library, and will be made available under an open-source licence. Detailed experiments were performed with the prototype to provide evidence of the mechanisms' efficiency and reliability even in the presence of faulty or malicious peers.

Bandwidth Trading on Demand

While PeerMart is generally applicable to any distributed trading scenario like P2P file sharing, Grid computing, or decentralized storage systems, the approach is particularly interesting for bandwidth-trading scenarios. Electronic marketplaces for trading bandwidth have emerged since the late 1990s, but were seriously hit by the economic downturn in 2001. However, driven by recent technical advances, especially in the area of optical fibre technology and virtual router infrastructures, the provision of bandwidth services 'on demand', being offered just like other commodities, will become a reality.

Today, bandwidth services are normally provided under the umbrella of long-term bilateral peering and transit agreements between individual Internet service providers and customers. Appropriate mechanisms for trading such services in an efficient and scalable fashion among multiple providers and customers are not yet available. A new workshop on bandwidth on demand (BoD) has therefore been established to bring together researchers from both industry and academia. The scope of the workshop includes the technical and economic dimensions of BoD mechanisms, legislative and regulatory issues, and the industrial development of new technology that supports these mechanisms.

To meet the individual requirements of bandwidth services, the generic PeerMart design is currently being refined and extended in a follow-up post-doc project at the University of Zurich, IFI. This includes the support of key service

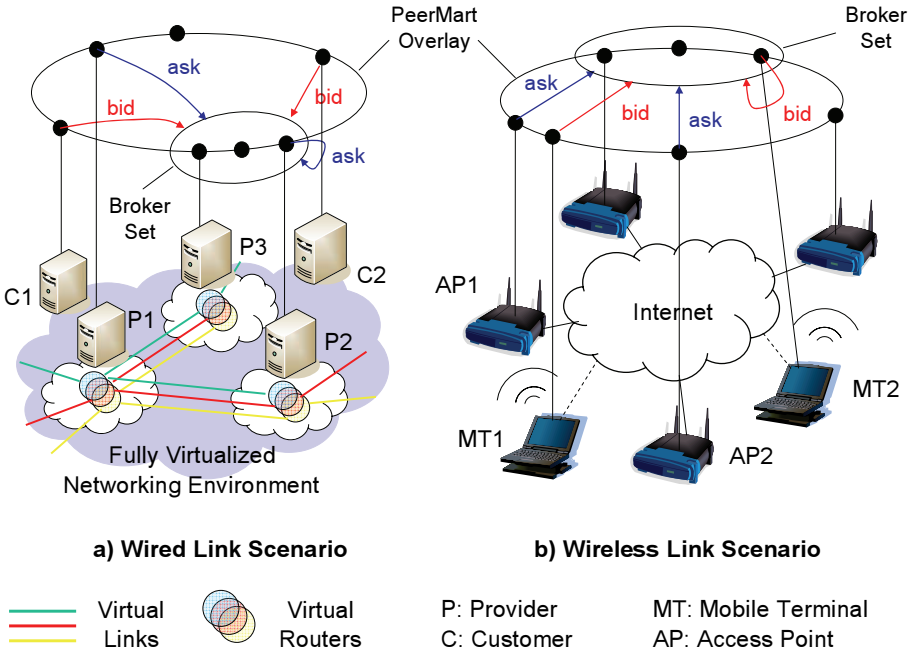

Figure 2: Bandwidth trading scenarios.

parameters like bandwidth, reliability, delay and jitter, as well as start-time and duration of such services. The adapted trading mechanism will enable users to buy and sell bandwidth services on demand or in advance, and will also allow the reselling of unused bandwidth services to other users.

Application Scenarios

As depicted in Figure 2, two different applications have been defined which will benefit from decentralized support for bandwidth services. One application is targeted at trading bandwidth over optical links in a wired, fully virtualized networking infrastructure. The other focuses on trading Internet services over access points between Internet access providers and end-users in wireless hotspots.

Fully virtual network environments are a promising future platform for Internet service providers and large end-customers, who may use them, for example, to sustain parts of an ISP's network or to carry traffic for broadcasting large sporting events. Using synchronous optical network (SONET) circuits, or wavelength-division multiplexing (WDM) technology supporting 'virtual' lightpaths, multiple bandwidth-guaranteed tunnels can be carried over a single optical dark fibre. In conjunction with new router platforms such as Cisco's CSR-1, which have the ability to provide secure virtualisation, routers themselves can be separated into multiple virtual routing engines. This enables the deployment of dedicated, bandwidth-guaranteed networks on top of fully virtual networking infrastructure. Using PeerMart's trading mechanisms, customers and providers of such a virtual network platform will be able to buy and sell fully virtual networks composed of virtual nodes and network links across different physical locations, adapting their capacity as necessary.

The second application shows that PeerMart can also be applied to wireless scenarios. At wireless hotspots like airports, train stations or other crowded areas, there are typically multiple users equipped with mobile terminals (laptops, mobile phones, PDAs etc) seeking access to the Internet. At the same time there may be several Internet access providers available and running multiple wireless access points. With PeerMart, it is possible to determine prices for such services in an economically efficient and scalable way based on the current supply and demand. Moreover, through the use of redundancy PeerMart can ensure high reliability of its trading mechanisms even in the presence of malicious or unreliable peers. Future work will include an integration of the developed bandwidth-trading mechanisms into real hardware devices such as open-source routers, access points, and end-user devices, which will enable a deployment of these mechanisms into production environments such as Internet eXchange Points (IXPs) or wireless hotspots.

Links:

PeerMart Project:
http://www.peermart.net/

PeerMart Dissertation:
http://hausheer.osola.com/publications/dissertation.pdf

BoD Workshop:
http://www.csg.unizh.ch/events/bod06/

FreePastry Project:
http://freepastry.org/

Please contact:
David Hausheer
University of Zurich, IFI, Switzerland
Tel: +41 44 635 43 72
E-mail: hausheer@ifi.unizh.ch

Net-WMS - A New Generation of Warehouse Management Systems Networked Services

by Francois Fages and Abder Aggoun

Net-WMS - Towards integrating Virtual Reality and optimisation techniques in a new generation of networked businesses in Warehouse Management Systems (WMS) under constraints - is a new project managed by ERCIM, developing interactive optimisation tools and prototype software that will form the basis for a new generation of networked services for WMS.

Net-WMS will handle networked communication and co-operation processes through the integration of decision-making technologies, generic 2D, 3D and higher-dimensional placement constraint solvers, visualisation and interaction with the solvers in virtual reality, packing models and knowledge modelling with business rules. Its scientific outcome will be relevant to the whole domain of combinatorial optimisation and will have direct technological impact on supply chain management at both the WMS and Transportation Management Software (TMS) levels, especially in the areas of packing, vehicle loading, space management, planning and scheduling, inventory control and packed item visualisation.

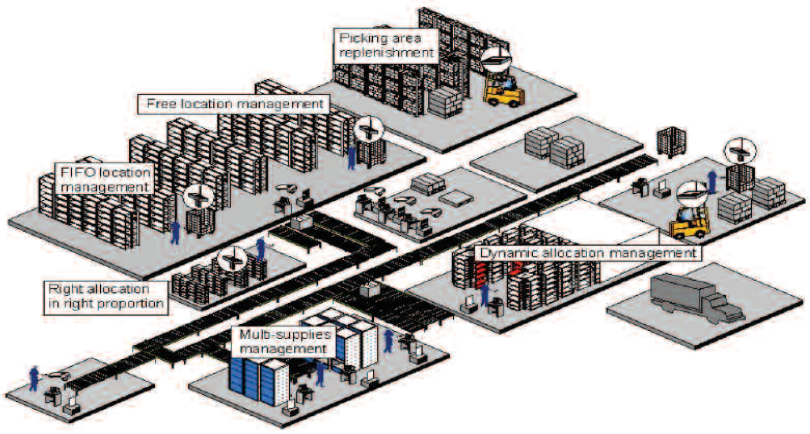

Figure 1: Example of an industrial warehouse powered by GILDAS WM, a warehouse management system from KLS Logistic System (Grenoble, France).

Prototype Applications

Net-WMS prototype applications will be operational in an environment of networked warehouses, thanks to their interoperability and to the user-friendly interface designed for plant-level technicians. They will include innovative tools and plug-ins including:

- a packing modeller of items based on optimisation techniques and interactions in virtual reality
- a palletizer tool using optimisation techniques
- a dispatcher, which includes the virtualisation of a truckload
- a set of interfaces enabling communications between several planning components across a network.

The project goals are threefold. First, from a scientific standpoint, significant advances are expected on:

- the algorithmic treatment of global placement constraints for objects of higher dimensions including space and time
- the expression of constraint optimisation problems with a language of business rules
- the control of an optimisation tool with interactions in virtual reality.

Second, on the technological side, the project will pave the way for next-gen-

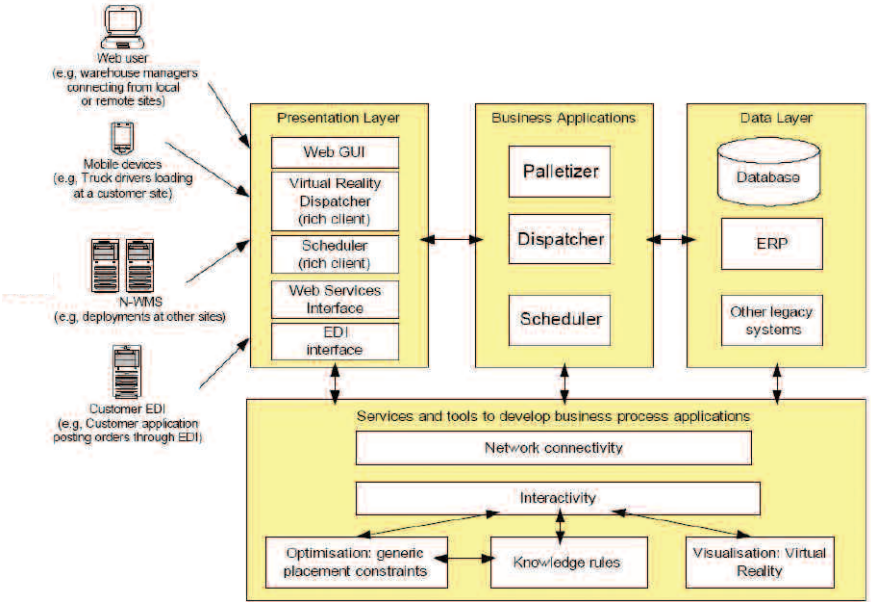

Figure 2:
Innovative Net-WMS J2EE architecture.
Applied technologies:
- Java Platform, Enterprise Edition (J2EE)
- Unified Modeling Language (UML)
- Model-View-Controller (MVC)
- Model Driven Architecture (MDA)
- Single Sign-On (SSO)
- 3D Visualisation.

eration WMS software by applying innovative technology to enhance operations in industrial warehouse environments. This includes:

- a set of Java Platform, Enterprise Edition (J2EE) interfaces for interoperability and mobile services, enabling communications between planning components across a network
- a mobility interface, allowing remote users (eg truck drivers) to report planning changes
- new interactive modules combining constraint programming, rule programming and virtual reality, in support of modelling, simulation and optimisation of the packing process
- a set of high-level modelling libraries for the constraint programming system Choco
- extensions to rule programming tools such as constraint handling rules (CHR) and Drools.

Third, on the commercial side, Net-WMS aims at improving European competitiveness in the area of warehouse management by significantly reducing costs related to packing, manpower and transportation.

Net-WMS is a Specific Targeted Research Project co-funded by the European Commission's ICT for Enterprise Networking D5 Unit. IST in FP6 focuses on future-generation technology in which computers and networks will be integrated into the everyday environment, allowing easy access to a multitude of applications and services through user-friendly human interfaces. Net-WMS places this vision at the heart of its activities and as such contributes to the Ambient Intelligence (AITPL) European cluster.

Net-WMS commenced on 1 September, 2006 and will be active for three years. The consortium has a combined expertise and field knowledge that guarantee the project will reach its objectives. It is composed of ten members representing both academia and industry including the ERCIM members INRIA and SICS.

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

Please contact:
François Fages
INRIA, France (scientific coordinator)
Tel: +33 1 39 63 57 09
E-mail: francois.fages@inria.fr

Abder Aggoun
KLS OPTIM, France
(technical coordinator)
Tel: +33 1 60 14 66 85
E-mail: abder.aggoun@klsoptim.com

Philippe Rohou
ERCIM office
(administrative contact)
Tel: +33 4 97 15 53 06
E-mail : philippe.rohou@ercim.org

The IBM Secure Trade Lane Solution

by Francois Dolivo

The IBM Secure Trade Lane (STL) is a new comprehensive global logistics information platform. It gives global supply-chain stakeholders access to information on demand, allowing real-time access and response to physical cargo monitoring data as well as the related logistics transaction data. For the first time, shipments can be monitored from the manufacturer to the store and related activities such as port operations optimized.

The international shipping industry plays a key part in the global economy. It is responsible for the carriage of 90% of world trade, with 50,000 merchant ships annually transporting over six billion tons of goods in some twenty million maritime containers. Without shipping, intercontinental trade, the transport of raw materials, manufactured goods and food would simply not be possible.

In this age of 'just-in-time' manufacturing and supply, containers are virtual warehouses that move goods from origin to destination via an intricate set of processes. These often involve more than ten different service providers, government representatives or intermediaries. Market and customer pressures demand that these complex operations be optimized, and this has implications for both public and private parties in the trade lane. Governments are therefore under pressure to improve clearance speeds and reduce costs. In addition, security must be drastically improved in this time of increased terrorist threats and security breaches: currently only 2-4 % of transported containers are physically inspected. These needs are driving changes in public policy that will affect the entire industry, with new challenges

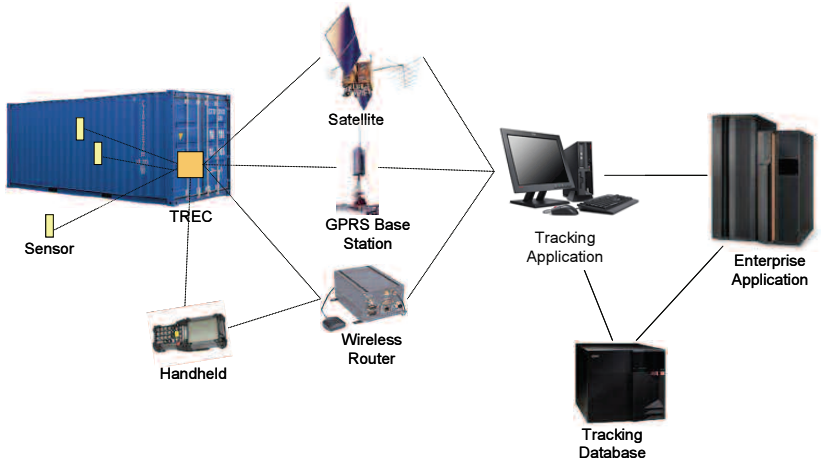

Figure 1: Overall view of the STL architecture.

arising in terms of security, reliability, liability, visibility, and efficiency of container shipment.

The IBM Secure Trade Lane (STL) is a new comprehensive global logistics information platform that addresses all these challenges by providing unprecedented levels of supply chain efficiency and security. It gives global supply-chain stakeholders access to information on demand, allowing real-time access and response to physical cargo monitoring data, and related logistics transaction data such as order information, invoices, financial data, bills of lading and manifests.

An overall view of the STL architecture is shown in Figure 1. At its heart is the TREC (Tamper-Resistant Embedded Controller), an intelligent wireless monitoring device (see Figure 2) that is mounted on the container. Information provided by TRECs is made available to supply chain participants through the Shipment Information System (SIS; see Figure 3). The SIS is a distributed network based on service-oriented architecture (SOA), which enables end-to-end data collection and reporting. It coordinates the sharing of information across authorized parties using proven techniques and tools.

The TREC platform has two primary functions. The first is to create an audit trail of container movements and events from the point of origin to the destination. The second is to make this information available to authorized entities, allowing them to perform risk analyses, to assess the container's security and integrity and to optimize the efficiency of container shipments. The TREC device automatically collects information on container events, including its physical location (based on GPS) and state (eg temperature, humidity, ambient light, acceleration and door status). It can communicate with the backend server via a satellite network, a cellular system (GSM/GPRS), or a Wireless Personal Area Network (WPAN) based on ZigBee/IEEE 802.15.4 radio. A handheld can also be used to communicate with the TREC over a WPAN. The TREC incorporates significant processing power, enabling it to analyse events and take appropriate actions. For example, opening of the door within a predefined geographic zone by an authorized person is an event that is logged; but opening of the door outside of the predefined zone or by an unauthorized person will trigger the sending of an alert to the backend server monitoring the status of the containers.

Figure 2: TREC prototype ready to be mounted on the container door. On the left is the box containing the electronics that sits within the container; at the bottom right is the antenna box that sits outside the container door.

As shown in Figure 3, the TREC devices are wirelessly connected to a Logistics Service Provider (LSP) through a secure, fully integrated network that links to the SIS. This system enables manufacturers, customers and any authorized member in the logistics network to check cargo integrity and location. Participants manage their own data through decentralized databases, ensuring full customer control over data access and privacy.

The IBM Secure Trade Lane solution will be composed of four service packages:

- Monitoring services enable complete real-time monitoring of any container, anywhere in the world. Collected data can be used to offer services to other trading partners who seek to better manage on-time deliveries, quickly identify bottlenecks, take advantage of green-lane customs treatment and realize additional supply chain efficiencies.
- Information-sharing services allow trading parties within and across industries to exchange data more easily, helping them in meeting their security and efficiency objectives.

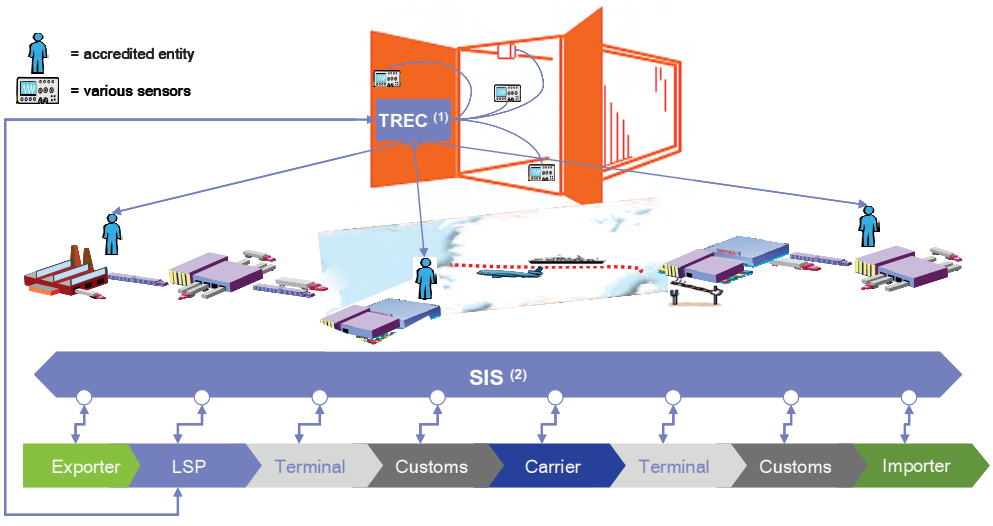

Figure 3: Shipment Information System (SIS) - an SOA-based distributed network enabling end-to-end data collection and reporting.

(1) TREC= Tamper-Resistant Embedded Controller
(2) CIS = Shipment Information System (Service oriented infrastructure)

- Efficiency services help exporters and importers to optimize their supply chains.
- Security services help any trading party to balance efficiency and security challenges in the context of their specific industry.

The STL concept and the first TREC prototypes were developed at the IBM Zurich Research Laboratory with the support of IBM Business Consulting Services. In September 2005, IBM and Maersk Logistics, of the A.P. Moller-Maersk Group, announced a partnership for pursuing the development of STL. Several hundred TREC prototypes have now been built by the IBM Engineering and Technology Services in Mainz. The SIS prototype is in development at the IBM European Business Solution Centre (EBSC), La Gaude, France, and the IBM Global Business Solution Center, Bangalore, India. IBM is participating in a large-scale international research project sponsored by the European Union, known as Information Technology for Analysis and Intelligent Design for e-Government (ITAIDE). The intent of the project is to define and, using STL, pilot ways to make international trade safer while reducing the administrative burden.

Links:
http://www-306.ibm.com/software/solutions/LE/LG04-01/solutions_overview.html

http://www-03.ibm.com/industries/government/doc/content/bin/IBM1092_03_ISO.pdf?g_type=pspot

Please contact:
Francois Dolivo
IBM Research GmbH, Zurich Research Laboratory, Switzerland
E-mail: fd@zurich.ibm.com

Clipped RFID Tags Protect Consumer Privacy

by Günter Karjoth and Paul Moskowitz

Existing methods designed to protect consumer privacy in RFID either put the burden on the consumer or suffer from the very limited capabilities of today's RFID tags. By using 'clipped' tags, consumers are able to physically separate a part of the antenna from the tag in an intuitive way. Such a separation provides visual confirmation that the tag has been turned into a proximity tag. Deliberate action on the part of the owner is then required to permit the RFID tag to be read. This mechanism enables controlled reuse after purchase, making the clipped tag a viable addition to the privacy mechanisms proposed for the use of RFID by consumers.

Radio Frequency Identification (RFID) tags typically are small devices that can be embedded in or attached to objects for the purpose of identifying the object over a radio channel. Objects tagged with RFID technology can be read more easily and more frequently, thus improving the quality of information on objects in a supply chain or in the inventory of a warehouse. RFID tags can be read if they are within range (typically up to a few metres) of a reader, which communicates with tags over a radio channel without requiring a line of sight. However, these characteristics of RFID tags have raised privacy concerns.

Retailers are constrained in the technology they use to protect consumer privacy. Stringent cost requirements limit the computational power of RFID tags, which in turn limits the mechanisms that could give users control over the use of their data in back-end systems. Existing solutions either put the burden on the consumer, which may include the risk of illegal behaviour, or are hampered by the very limited capabilities of inexpensive tags. Until now the 'kill' command seems to be the solution with the greatest potential. However, it is still necessary to overcome its three major weaknesses: complex key management, no (controlled) reuse after purchase, and no (visual) confirmation of successful disablement. In response to these limitations, we propose an RFID tag structure that permits consumers to disable a tag by mechanically altering it. We call such structures 'clipped tags', since a portion of the antenna is separated from the tag.

The mechanical alteration inhibits the ability of a base station or reader to interrogate the RFID tag or transponder by wireless means, and provides visual confirmation that the tag has been deactivated. Once a tag has been disabled (or 'clipped'), the distance over which it can be read is drastically reduced. Later use would require deliberate actions on the part of the owner of the RFID tag to permit the tag to be read, and thus could not be undertaken without the owner's knowledge unless the item were either stolen or left unattended. This makes it an appropriate mechanism to implement consumer consent.

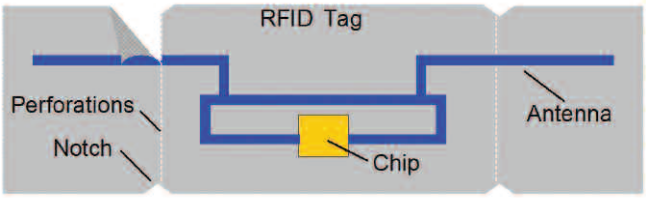

Figure 1: Schematic scheme.

Clipped tags are a simple and practical privacy-enhancing technique for RFID retail. Disabling can be performed in an easy, reliable and verifiable way. Even if the RFID tag is 'printed' onto a prod-

uct, its antenna can be partially removed. Whereas the physical destruction of a tag would likely damage the original item, clipped tags provide a practical way to separate the antenna. For instance, we have demonstrated paper labels with an RFID tag inside. The tag comprises an antenna, which is printed, etched or stamped onto a substrate (eg a plastic foil), and a silicon chip attached to it. If necessary, another plastic foil may cover the tag to protect it from inclement environments. Such labels are then affixed to objects, and stored information may be written and rewritten to the embedded chip.

There are a number of possible ways of making clipped tags. One preferred structure uses perforations like those on postage stamps, which are manufactured into the antenna and its substrate. A separation along the line of small holes or cuts then detaches a portion of the antenna from the tag. Figure 1 shows a schematic of a clipped tag employing perforation. The notches additionally provide tear initiation, similar to the slits placed in ketchup packets. When torn, a portion of the antenna is removed.

In another example, the antenna may be constructed of conducting 'scratch-off material'. This material is familiar to consumers from its use to obscure printed material on lottery tickets or prepaid phone cards. The antenna of the RFID tag is manufactured on a substrate using the scratch-off material.

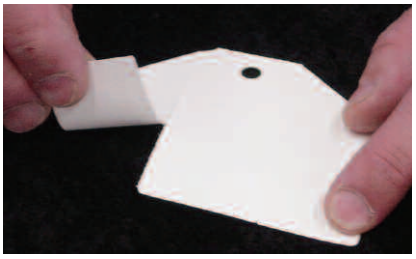

Figure 2: Partial antenna removal from a garment hang tag.

The substrate or mount may be a plastic material such as polyimide or polyester. The chip is mounted on the substrate and is connected to the antenna by an electrical conductor or conductors. The RFID tag is manufactured in such a way that a part or all of the antenna or its connecting wiring is exposed. The electrical conductor or conductors pass through a window, such as an exterior portion of the substrate or mount. Such tags are placed on the article or on its packaging in such a way that the antenna or the antenna-chip connection can be scratched off using a coin or fingernail. The consumer or a check-out attendant in a retail establishment may then perform the scratch-off operation to disable interrogation of the tag. The tag is open for visual confirmation that the tag has been deactivated.

To demonstrate that such physical structures can be embedded in today's manufacturing process, IBM has been collaborating with Marnlen RFiD, a leading manufacturer of RFID labels, to design and manufacture samples of garment hang tags. The starting point was the Alien Technology, ALL-9440 (Squiggle 2.2) Gen2 UHF RFID tag inlay. This tag may be read at distances of 30 feet (10 metres). The Alien tag inlay was incorporated into a printable garment hang tag or label by Marnlen RFiD. We have also experimented successfully with inlays from other manufacturers: Symbol, Intermec, and Impinj.

The label contains notches for tear initiation and perforations to direct the tear. The design is consistent with the printing of graphics and the writing of tag information in a standard RFID printer. Marnlen RFiD manufactures clipped tag labels using standard RFID inlays. The tag design is compatible with commercially available readers and printers. Figure 2 shows an RFID tag inlay incorporated into a printable garment hang tag or label. The hang tags may be torn at the point-of-sale by the consumer to remove parts of the antenna. The result is a reduction in tag read range from 10 metres to about 2 cm.

Link:
http://en.wikipedia.org/wiki/Clipped_Tag

Please contact:
Günter Karjoth, IBM Zurich Research Laboratory, Switzerland
E-mail: gka@zurich.ibm.com

Paul Moskowitz, IBM Watson Research Center, Hawthorne, NY, USA
E-mail: mosk@us.ibm.com

TraSer – Identity-Based Tracking and Web Services for SMEs

by Zsolt Kemény and Marcell Szathmári

Today's trends in industrial production are marked by heterogeneous conglomerates such as production networks, and rapidly changing customer requirements. In order to improve competitiveness in such a production scene, tracking and tracing of both goods and data must remain efficient across company borders. This is, however, associated with a high investment burden, given today's technological background. The EU-funded project 'TraSer' aims at reducing these costs by providing free, open-source solutions for tracking and tracing across company borders.

In responding to more intricate customer demands, today's industrial production has to deliver more than a mere product, since in numerous cases, customers wish to receive up-to-date information about the status of their individual orders. Moreover, the custom-tailored production of items is spreading, even in areas where simple 'make-to-stock' policies used to prevail in the past. All these trends necessitate the improvement or introduction of item-level support for tracking (what is the status of the item itself, eg where it is located during delivery) and tracing

(what happens to the item with respect to production processes, eg what other product is it built into). Given today's heterogeneous production networks, these solutions often have to reach beyond company borders while still providing services in a uniform way. While technically feasible solutions have emerged in the past, in most cases they are suitable only for large companies, and small and medium-sized enterprises (SMEs) still find it difficult to join such tracking and tracing efforts, even though participating in them could be vital for groups of cooperating small companies.

In order to overcome these difficulties by providing a more SME-friendly solution platform, the international research project 'TraSer' (Identity-Based Tracking and Web-Services for SMEs) was started in June 2006. The efforts are planned to culminate in the creation of an open-source TraSer community. Through its Web services, this will assist the targeted group of users (mainly SMEs) to improve their efficiency while operating in changing environments such as supply-chains, service operations and project delivery networks. The main output of the project, the TraSer open-source platform, is based on experience gained during earlier work at the Helsinki University of Technology, the development of the DIALOG system (see the DIALOG Project Website, http://dialog.hut.fi). The expected benefits are as follows:

- A low initial investment will be needed, as users will be able to set up a basic product-centric tracking service with components available for free at the open-source community.
- Lean implementation and maintenance will be possible, as the application of the solution platform is meant to be simple, not requiring many IT specialist staff.
- By allowing the use of the 'ID@URI' item identification scheme, as already employed in the DIALOG project, it will be possible to spare the costs of ID registration, since – in contrast to present-day proprietary ID providers – the participants can manage distributed unique ID allocation on their own. In addition, up-to-date status information on the items will be stored by the manufacturer in charge (as opposed to paying for centralized tracking services provided by third parties), making a further reduction of operating costs possible.
- The TraSer solution platform is intended to guarantee maximal compatibility with legacy and low-end standard systems, as the community will provide the tools for embedding new tracking and tracing services into existing software environments, such as enterprise resource planning (ERP). It will also be possible to adapt the item identification scheme to legacy systems, meaning it will become feasible to integrate traditional numbering schemes already in use by one or more partners. As a further concession to 'as-is' situations, detailed and freely available documentation will allow companies to re-implement parts of the solution platform within a software environment of their choice if they wish to reduce the diversity of software technologies used in their own system.
- Companies using TraSer will have easier access to current logistics data and the RFID and barcode readers of service providers, as TraSer will deliver guidelines and standards which, when followed, will ensure interchangeability with information from other SMEs and logistics service providers (LSPs).

In addition, the following outputs are planned for the project:

- industrial implementations
- reports on the implementation cases
- descriptions of specific SME business scenarios, and models of the potential benefits of developing product-centric applications in these scenarios
- advancing theory on product-centric service innovation, the evolution paths for product-centric information architectures and applications.

This project has been funded for three years by the 6th Framework Programme of the EU. It is a collaboration between SZTAKI (project leader), the Helsinki University of Technology and the University of Groningen (RuG) representing the academic sector, with industrial partners Innotec Magyar Kft., Finland Post Corporation, Netherlands Organisation for Applied Scientific Research (TNO) and Wittmann & Partners Computer Systems. These commercial organisations share practical background knowledge and provide a testing terrain for the first application cases of TraSer. Further information about the project can be obtained at the official TraSer website or by contacting the representatives of the participating organisations.

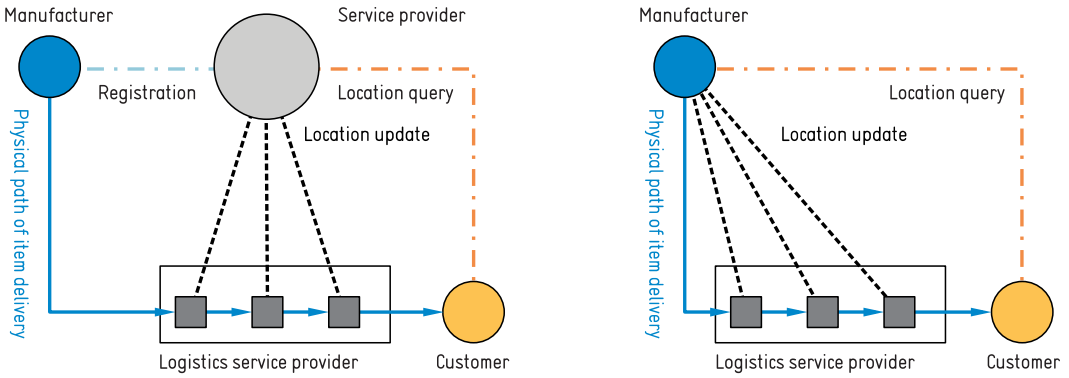

Tracking during item delivery using third party proprietary services (conventional approach, left) and using decentralized item numbering and tracking with the TraSer solution platform (right). The latter scheme is expected to spare considerable setup and operating costs.

Link:
http://www.traser-project.eu/

Please contact:
Elisabeth Ilie-Zudor and Zsolt Kemény
SZTAKI, Hungary
Tel: +36 1 279 6195
E-mail: ilie@sztaki.hu,
kemeny@sztaki.hu

Distributed Engine for Advanced Logistics

by Tamás Máhr and Alfons Salden

How can agent-based technology reduce transport costs, traffic jams and carbon dioxide emissions? DEAL, a Dutch research project on multi-agent-based logistic monitoring and planning, is studying these questions.

The research and industry partners in the Distributed Engine for Advanced Logistics (DEAL) project are Almende BV, RSM Erasmus University, Vrije Universiteit Amsterdam, CWI, CarrierWeb, Post Kogeko and Vos Logistics. The goal of this consortium is to apply new knowledge to a broad range of products for national and international logistic services. In particular, it aims to speed up negotiations amongst software agents representing different stakeholders in dynamic logistic problems. This would enable an increasing number of logistic conditions and alternative logistic plans to be examined, and therewith a higher overall optimal performance of the logistic system to be achieved.

agent-based technologies can help reducing transport costs, traffic jams and carbon dioxid emissions .

The current best practice for planning in a logistics company consists of human planners interacting with an order-assignment database and a tool for calculating optimal routes. The most advanced logistics tools focus on balancing well-defined goals such as minimizing the number of vehicles, the distance covered, and the number of empty kilometres driven. The task of the human planners is to make choices, often by imposing additional constraints on the optimisation tool, such that all involved parties are more or less satisfied with the schedule. This process is time-consuming, and a schedule may still be sub-optimal when the time comes for it to be executed.

Another issue is that the extra constraints imposed by human planners on the optimisation tool consist of a large number of rules that express the preferences of all the involved parties. To learn all the rules is a lengthy process and makes the logistics companies vulnerable to changes in personnel.

Outperforming the Optimum

The DEAL project develops a multi-agent based logistic planner and monitoring tool that allows planning and execution to overlap. Truck and container agents dynamically change plans to react to unexpected events by using different market mechanisms such as auctioning and de-commitment. The agents continuously search for better solutions while always maintaining an executable plan. Even while some parts of a plan are being executed, the agents keep searching for better solutions for the rest of the work, and if any incidents occur, they adjust their plans immediately.

Another issue studied in the DEAL project is the evaluation of logistics performance. It is argued that simply evaluating one of the traditional performance indicators is insufficient for a modern logistics provider. Rather, monitoring of key performance indicators of all the stakeholders is essential for successful business relations. This means that the optimisation goals of the agents are manifold. To decrease the number of optimisation goals, we introduced an evaluation framework based on fuzzy logic. This translates the values of several key performance indicators into satisfaction measures of the main stakeholder groups: management, customers, employees and society. Given these measures, agents can concentrate on generating a solution that balances the satisfaction of these four groups.

In conclusion, we see the problem faced by logistics companies to be not merely an optimisation for cost problem, but a balancing problem within an uncertain and unstable environment. Companies must provide plans for transportation vehicles such that all parties are satisfied (including management investments), while constantly facing incidents that force changes of plan.

Link:
http://www.almende.com/deal/demo

Please contact:
Tamás Máhr and Alfons Salden
Almende, Rotterdam, The Netherlands
Tel: + 31 10 4049444
E-mail: tamas@almende.com, alfons@almende.com

Supporting City Planning with Traffic Monitoring in Image Sequences

by Matthieu Molinier

Traffic monitoring is becoming increasingly important in the planning of transportation and infrastructure in urban areas. Video sensors can capture detailed information about the nature and speed of moving traffic objects. VTT has developed a generic tool for traffic monitoring, which allows the interactive definition of traffic events and stores position and speed records in geocoded (GIS) form for later analysis. This tool is being used by the Helsinki City Planning Department for intersection monitoring and for the joint analysis of vehicle and pedestrian behaviour near pedestrian crossings.

As traffic volumes increase in urban areas, so does the need for and interest in traffic monitoring for purposes such as city transportation planning or safety analyses. With the improvement of image-processing techniques, video sensors are becoming one of the natural successors to loop-detectors for traffic studies. Video allows traffic objects like buses, cars, pedestrians and bicyclists to be identified, and their speeds estimated. On a higher level, video analysis can help in describing the interactions between traffic objects and in detecting events or incidents such as failure to give way. Yet it remains difficult to provide reliable tracking in congested conditions that generate highly occluded scenes. VTT is developing a generic tool for traffic monitoring that can extract relevant traffic parameters from videos and store the results in GIS (Geographic Information System).

Method development started on image sequences acquired from helicopters, as these can cover a greater area than ground-based systems. The methods were then applied to static ground-based cameras, on which the majority of traffic-monitoring systems are based. The processing chain includes image geocoding and moving-pixel detection followed by tracking of moving objects. The first image of the sequence was manually registered to a GIS map to allow vehicle speeds to be later evaluated in meaningful units.

The pixels in motion were then detected by adaptive background subtraction, which iteratively constructs an image of the background (static pixels), and subtracts it from the current image in the video. Moving objects were aggregated and tracked throughout the sequence by spatiotemporal connected-component analysis. This simple approach allowed

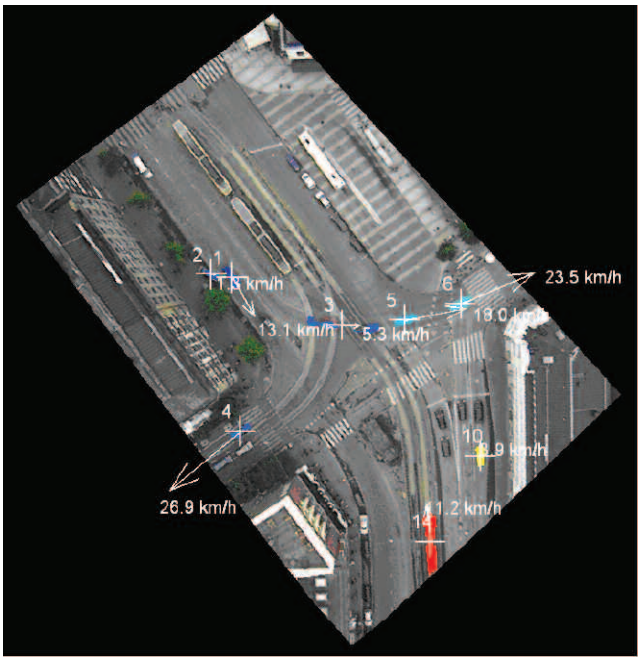

Figure 1: Vehicle-speed estimation in an image sequence acquired in near-nadir direction from a helicopter flying at an altitude of 200 metres over an intersection in Helsinki.

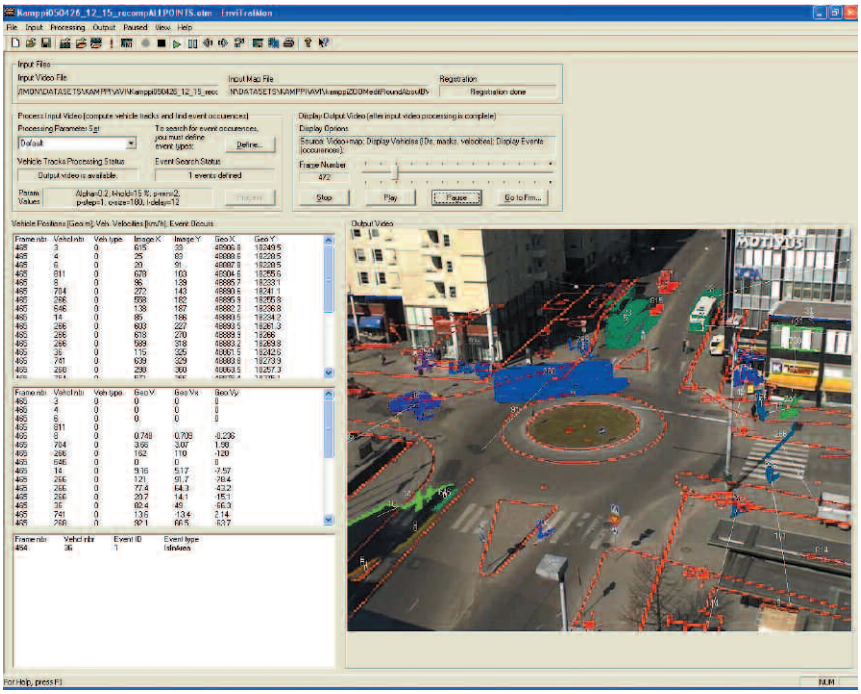

Figure 2: TRAFMON, a tool for interactive traffic monitoring. Red lines symbolize the GIS map overlaid on input imagery. In replay mode, the sequence can be browsed in search of specific events. At every frame of the sequence, position and estimated vehicle speeds (in km/h) are displayed on the left-hand side.

most of the noisy pixels to be discarded while spatiotemporally consistent targets were tracked. In case of helicopter-based sequences, semi-automatic image-sequence stabilisation was carried out as pre-processing. The methods were successfully tested in an aerial image sequence acquired at an altitude of 200 metres (Figure 1), as well as with ground-based cameras.

The methods have been integrated with a piece of software known as TRAFMON, which has a Graphical User Interface (GUI) and allows semi-automated and interactive analysis of image sequences for traffic monitoring. The user can draw a zone of interest in the scene and associate events with it, eg a vehicle stopping or slowing down in this zone. The 'philosophy' of the tool is to aid later analysis by extracting from every frame of the video all the relevant traffic information (eg vehicle positions and speeds in real-world coordinates, and triggered events). Records are then stored in tabulated text files, and can be used in replay mode for displaying information (Figure 2). They could easily be used for computing traffic density, vehicle speed profiles or other parameters useful to traffic engineers. The huge amount of data generated by several minutes of video can be used in data mining; for example, identifying typical vehicle behaviour at a given intersection. The tool has proved useful at Helsinki City Planning department, assisting intersection-monitoring tasks that were previously done manually.

This work was part of a larger project, ENVIMON (Environment Monitoring using Earth Observation), the purpose of which was to develop a general software platform utilizing Earth Observation data. ENVIMON benefits the TRAFMON software by providing a set of tools for image pre-processing and post-processing; where necessary, these can also enable the distribution of data and processing results via the Internet. Further work in the traffic-monitoring project will be supported by the Finnish Ministry of Trade and Industry and the Helsinki City Planning Department, and will include making improvements to the tracking procedure for ground-based imagery and to the handling of occlusions of vehicles by other vehicles. The ability of TRAFMON to define events involving more than one object, eg two vehicles approaching each other in a collision course, will be enhanced.

We are looking forward to developing collaborations with traffic consulting companies, and possibly with other European research institutes within the FP7 program – partners are welcome!

Link:
ENVIMON Traffic monitoring website:
http://virtual.vtt.fi/virtual/proj2/envimon/traffic.htm

Please contact:
Matthieu Molinier
VTT, Finland
Tel: +358 20 722 7345
E-mail: matthieu.molinier@vtt.fi

Traffic Planning and Logistics: A View from the Edge

by Antonio Lucas

Past and present visions of an effectively managed and integrated road transport system have in general missed a key point: the informed and cooperative road user. This is disappointing, because even in the most sophisticated scenario (eg ADAS), the power of decision and the freedom to use electronic devices will be on the driver's side. Comprehensibility and credibility of devices within the ITS domain are essential for road users and thus for the improvement of the entire transport system. Recent European initiatives both within DG TREN (Easy way) and DG INFSO (e-Safety) have brought new attention to this important issue.

In 1970 more than 90,000 drivers died on European roads, and congestion and accidents stood out as the main issues to tackle. Within the COST framework (European Cooperation in the field of Science and Technology), 'electronic aids' helped to optimize existing transport networks through the considered use of road information. In the 1980s and 1990s, radio broadcasting and variable message signs (VMS) emerged as the first systems to bring real-time information to the road environment.

Since then, a variety of bodies have worked to develop a specific context for telematic implementation. Various projects exist within the European Framework Programs (Drive I and II, Transport etc), and the Directorate-General for Energy and Transport (DG TREN) has strongly promoted Euro-regional projects (ARTS, SERTI, CORVETTE etc) for Intelligent Transport Systems (ITS). In the last thirty years, these research efforts and road safety policies have contributed to the number of people dying annually on European roads being halved to 45,000. In 2001 the European Commission stated in the so-called Lisbon agenda that this death toll was still unacceptably high, and the challenge was then set to achieve a further 50% reduction by 2010. In achieving this, technology will again play a crucial part.

As a complex multinational and multilingual environment, Europe must work hard to keep up with ITS (now also e-Safe) technology as it is applied to road mobility and safety. Sadly, the European road user, the last link in the implementation chain, has somehow been neglected.

For instance, while VMS applications are now a mature technology (ie efficient and cost effective), the United Nations' Convention on Road Signs and Signals, which claims to have produced a catalogue of standard signs with which to inform road users worldwide, includes a bare minimum of useful pictograms to indicate variable or dynamic

events (eg pictograms for fog or 'ghost drivers' are still missing).

As far back as the 1970s, VMS were recognized as tools that would improve mobility and safety by virtue of traffic-flow re-distribution – thereby aiding traffic planning and logistics. However, after twenty years of VMS implementation, pictograms or graphics able to indicate or recommend rerouting schemes are only just beginning to emerge in Europe (see figure; to date only the congestion pictogram enjoys official acceptance in the UN catalogue).

Turning to more recent tools, the implementation of RDS-TMC (Radio Message System-Traffic Message Channel), which brings real-time traffic information to vehicles, has been affected by commercial and institutional procrastination. While countries such as Germany and Spain make traffic information freely available via public radio, some popular navigators are sold without TMC receptors. These navigators are therefore unable to take public traffic information into account, something many road users are unaware of. It appears that highly speculative 'business case' reasons are responsible for this absurd state of affairs.

Another inconsistency occurs in digital maps and navigators that do not include a subject-profile facility. For example, some tunnels in Spain have vehicle weight restrictions. Tunnel operators sometimes report trucks having to turn back and take an alternative route (specified several kilometres before the tunnel). The reason is simple: many truck drivers rely on navigators built for cars. These navigators thus neglect the obvious needs of a very important group, despite the growing volume of international freight and the consequences of this on planning and logistics.

VMS and RDS-TMC are 'old' technology, which after twenty years should have produced a greater benefit for society. Both DG TREN (promoting initiatives such as the Mare Nostrum project for VMS harmonisation in Europe) and DG INFSO (promoting TMC and user-awareness working groups within the broader eSafety initiative) appear to recognize this.

However, something is clearly missing from our modus operandi when it comes to planning the penetration of technology: either we forget about end users or we take them for granted. We live in an era that burgeons with technological innovations marked by unending acronyms: for example, ESP (Electronic Stability Program), V2V (Vehicle-to-Vehicle communication), V2i (Vehicle to Infrastructure), FCD (Floating Vehicle Data) and ADAS (Advanced Driver Assistance Systems). How will these innovations complement each other? What response is expected from users, and can we predict their behavioural adaptations? And most importantly, why is it still so common for highly sophisticated technologies to be designed, developed and promoted "for the road user but without the road user"?

On top of this, road users are clearly paying the toll of inconvenient events and avoidable road accidents. The introduction into the market of new navigation and road safety technology tends to be much slower and more difficult than for engines and mechanical devices. This is simply a failure of education: the user is unaware of or uninterested by them and as a consequence, many potential success stories have yielded to a massive reliance on hands, wheel, horn and blinkers. In the meantime, congestion and the optimisation of road flows and road infrastructures is still the main issue for our transport systems and will continue to be so for the foreseeable future.

PICTOGRAM	Developed or tested in	Status in UN
Congestion	1985 (COST30 bis)	Included, in use
Re-routing (several alternatives)	1985 (COST30 bis)	Just informed
	1990s (CENTRICO)	Unaware
Road closed, exit available	2006 (SOMS/IN-SAFETY)	Just informed
Exit closed (proceed)	2006 (SOMS/IN-SAFETY)	Just informed
Exit closed, alternative exit	2006 (SOMS/IN-SAFETY)	Unaware
High Occupancy Vehicle lane available	2006 (DGT/UVEG)	Just informed

Pictograms relating to traffic planning and logistics.

Links:
Mare Nostrum VMS project:
http://www.arts-mip.com/vms/default.asp

UN-ECE "Small Group on VMS harmonisation":
http://www.unece.org/trans/main/

SOMS/IN-SAFETY project:
http://www.insafety-eu.org/

eSafety initiative:
http://www.escope.info/

Euro-regional projects:
http://www.i2tern.info/

Please contact:
Antonio Lucas
Road Safety and Traffic Institute,
University of Valencia, Spain.
Tel: +34 96 3393880
E-mail: antonio.lucas-alba@uv.es

Software Travels in the Fast Lane: Good News or Bad?

by Fabrizio Fabbrini, Mario Fusani and Giuseppe Lami

History often repeats itself. In just a few years, the technical solutions for control software in the automobile industry have run through almost all the known stages of computer systems evolution. Progress has been extremely fast and the most advanced results in Software Engineering (SWE) and Telecommunications (TLC) are now finding their preferred deployment scenario within the automobile. This offers new services both to the driver and the passenger - but how much of this is good news?

In the near future, engine and passenger comfort control via x-by-wire, and time-constrained Web services, will no longer be optional or futuristic features. It is likely that the pressure of market competition in the automobile industry will compel both car makers and device suppliers to exceed even the most daring research challenges.

The positive side of this market-driven technological push is twofold. On the one hand, research is stimulated by continuous requests from the field. On the other, cars are becoming more comfortable, easier both to drive and to inhabit. They are more environmentally sustainable and are even better at fulfilling their primary function, ie transporting driver and passenger from A to B in the best possible way.

The System and Software Evaluation Centre (SSEC) at ISTI-CNR has been dealing with the effects of the increasing importance of automobile software (expected to reach 85% of car project budget by 2008). European car makers have been considering the relevance of the suppliers' software process since 2000. In 2001, SSEC convinced the main national Italian producer, Fiat Auto, to launch a program for supplier evaluation. Between 2001 and 2006, more than 25 suppliers from seven European and North American countries were process-assessed. We believe that most suppliers, not just Fiat Auto, have benefited from this.

However, progress can also bring problems. Even if 'problems' in a research environment may mean good news, we are concerned about some possible ways of dealing with innovation. Of course, it is too early to see whether features such as infotainment, city net services or Web services also imply a level of risk. Yet even though the driver and passengers in a car may be well aware of any in-car services (Web or no Web) and their qualities, software often acts on crucial vehicle functions in ways that are totally opaque, unforeseeable from the user's perspective and frequently purposely masked. For instance, the act of touching a lever to flash the lights no longer operates a switch, but signals the kernel of the operating system to schedule a predefined task that, when executed, operates the lights. This means that software is mediating the car controls. The popularity of 'x-by-wire' projects (x being a placeholder for 'brake', 'steer' etc) seems to indicate that this is going to be a prominent feature in the near future.

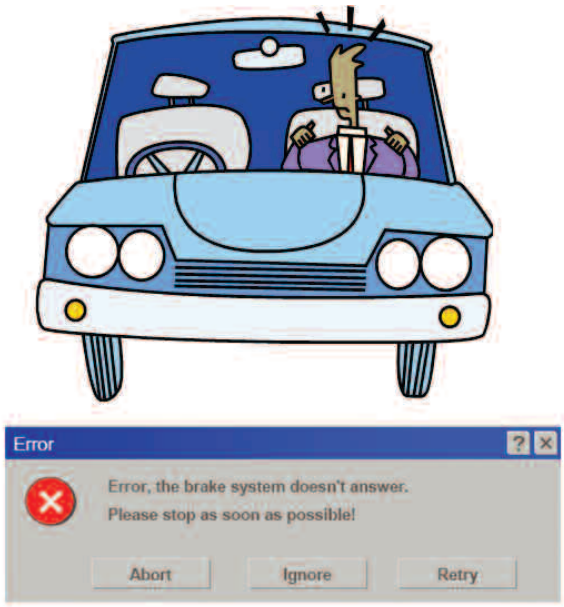

Software often acts on crucial vehicle functions.

Of course, users do not need to know how car controls are working, provided that the perceived performance is good. X-by-wire solutions were introduced to replace unreliable (and unsafe) mechanical devices and to perform complex functions like ensuring vehicle stability control in particular situations (eg acceleration, jerks, turns, braking). We know that part of this technology, the electronic components, can be highly reliable and safe. But we also know that software can fail.

Even though many believe that SWE is sufficiently mature to guarantee the reliability of software-controlled systems and that they meet safety requirements, what is not guaranteed is that car manufacturers actually adopt the most suitable SWE techniques.

Optimistically, there are standards that deal with these issues. One important standard is IEC 61508. This is a generic standard that has given rise to the emergence of specific standards in various application domains, such as train control and medical equipment control. It

has also generated a standard (still under development) that should cover safety in the automobile: ISO 26262.

However, we strongly believe that highly innovative solutions where software has direct control over critical car functions, such as the case of x-by-wire, should only be adopted with utmost care. Our concerns are based on several aspects, and in particular:

- in general, the standard ISO 26262 resembles part 6 of IEC 61508. This mainly concerns reliability, but has little to do with safety
- both IEC 61508 and ISO 26262 contain only a vague notion of process (they refer to the more traditional view of phased projects). We know from our experience how important it is that mature processes are executed by car makers' suppliers
- most supplier organisations that we have assessed (and most organisations assessed by other institutions) do not have sufficient process maturity to handle safety-critical projects
- safety is a system property, and can be understood and implemented having the vision of a system in which all the component interactions are well known (software is usually the weakest of such components). Software engineers alone cannot ensure safety; for example, in avionics safety is implemented at system level (and also regulated by different standards)
- when a standard is officially recognized by all the stakeholders, what then matters to the manufacturer is standard compliance. However, such compliance may not guarantee sufficient confidence with respect to 'real' system safety.

It is possible that not all suppliers or car manufacturers are in agreement with the above points. However, we hope that we have encouraged further efforts in software and system engineering with the objective of guaranteeing both the reliability and the safety of the diverse, complex and evolving field of automobile applications.

Please contact:
Mario Fusani
ISTI-CNR, Italy
Tel: +39 050 3152916
E-mail: mario.fusani@isti.cnr.it

Coping with System Complexity: Identifying Dichotomic Architectural Alternatives

by Gerhard Chroust

Certain types of information systems tackled today (so-called 'wicked systems') mean it is necessary to make some initial assumptions about the system architecture before even starting with the actual conceptualization. In general these assumptions define the architecture of the system only on a very high level, but with today's time and money restrictions, it is generally unfeasible or too costly to later change them. These assumptions can be classified into different dimensions using the concept of 'dichotomic architectural alternatives'. Here we discuss implications for the resulting system properties.

Today's information systems show a continuous growth in complexity. This growth was identified by Manny M. Lehman back in 1985, who placed systems into certain classifications. These were 'specification systems' (where a complete specification of the system is laid down and must then be fulfilled by the implementation, eg assigning carplates); 'problem systems' (where only a limited, incomplete knowledge of the problem and the resulting specification is available and the implementation is challenged to find acceptable solutions, eg a traffic simulation program); and 'environment systems' (which have the same properties as problem systems with the added difficulty that their very introduction into the system environment will change the problem, eg control of traffic lights based on simulation with the burden of the uncertainty of drivers' future behaviour). Hermann Kopetz has since added a fourth category – 'wicked problems' – which have the additional property of being large, complex, ill-defined and lacking a clearly identified objective (eg asking for the traffic control to also optimize the travel times of cars). He observes that such a system cannot be specified without some prior concept of its solution. Changing these concepts at a later stage cannot be done without considerable loss of time and effort. As a consequence, ill-conceived projects must either be continued with an inappropriate basic architecture or be abolished.

Classification of Architectural Alternatives

We observe that these basic concepts can be classified as 'dichotomic architectural alternatives' along different dimensions, usually providing an either/or alternative. In the real implementation a certain trade-off can usually be achieved, but the basic concept will stay. For example, should a system be centralized or decentralized? In everyday life, the proverb "You can't have your cake and eat it too" can also be seen as an example.

We intend to isolate and describe such alternatives together with their essential properties, implications, consequences and interactions. In this way, system designers (both newcomers and experts) will be given a clear understanding of the different options available to them, have some support for their intuition, and will therefore make better initial architectural decisions.

We classified the alternatives along the following dimensions:

- *Enactment time:* WHEN should a foreseen action be performed? (Example: prefetching of variables or just-in-time fetching)
- *Physical Location:* where should necessary data and programs be placed? (Example: centralized versus decentralized storage)

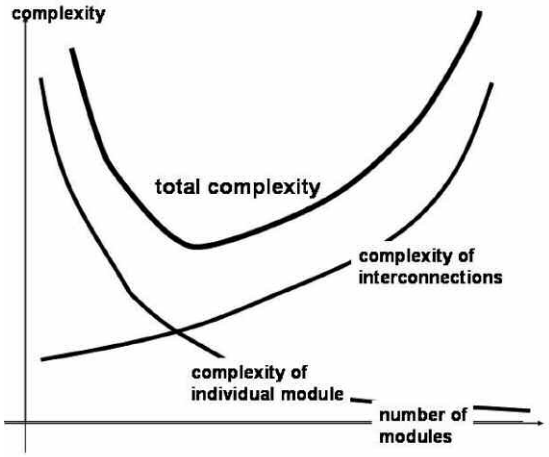

The trade-off between module size and number of interconnections, and their impact on system complexity.

- *Granularity:* how many/much should be accessed or handled in one step? (Example: choosing a large or a small page size for a paging system)
- *Communication Control/Responsibility:* who or what has the responsibility for leading the interaction? (Example: controlling the communication between sender and receiver either by polling or by an interrupt mechanism)
- *Risk expectancy:* how probable is an event? (Example: trying to prevent an (improbable) error versus providing detection and repair mechanisms)
- *Optimism versus pessimism:* should one prepare for the worst case or for the 'sunshine situation'? (Example: does one expect many or few errors in user inputs?)
- *Reductionistic versus holistic approach:* Should one consider only the problem at hand or all its ramifications? (Example: just planning a new motorway technologically or also taking into account environment and commercial implications)
- *Planning horizon:* how far into the future should the considerations go (Example: looking for short-term return on investments or for sustainability?)

Implication on System Characteristics

The end-user of a system is not really interested in the architectural choices but rather in the effects of those choices on system properties, such as runtime, size, complexity and development risk. (A more detailed list of system characteristics can be found in ISO/IEC 9126 and in effort estimation methods like COCOMO.) The next step is therefore to identify the implications of the architectural alternatives on these characteristics (for instance, decentralization is less vulnerable to attacks but more complex to implement).

Another issue is the cross-impact of these alternatives on one another and thus on the aggregated system characteristics. A typical example is the cross-impact of alternatives for choosing a module size: choosing a small module size (granularity) reduces the complexity of the individual modules but at the price of increasing the number of interfaces, which boosts complexity of the whole system. A compromise between the two will give some optimality (see figure).

Research in this area clearly holds many challenges. A better understanding of the alternatives and their effects will be of help to both novices and seasoned designers, and will also give a chance to teach in a holistic way some of the basic design decisions that must be made during the initial phase of a new challenging project.

Link:
http://www.sea.uni-linz.ac.at

Please contact:
Gerhard Chroust
J. Kepler University Linz / AARIT
Tel: +43 70 2468 8866
E-mail: gc@sea.uni-linz.ac.at

Power on Tap

by Lars Rasmusson

What would computing be like if we had a virtually infinite pool of computing power at our fingertips? And how would we go about creating such a pool? These and other related questions are the subject of research at SICS.

SICS is building infrastructure that will provide computing resources on demand, so-called 'utility computing' or computing power on tap. The name stems from an analogy made with traditional utilities, such as water and gas. Traditional utility resources are acquired as they are needed, without the need to reserve them in advance, and are paid for in proportion to usage. This is very different from the way in which computing power is acquired. We must buy the machinery that generates the computing power, rather than being able to easily buy the power only.

Who will benefit from having computing power on tap? The initial beneficiaries will be companies with fluctuating workloads and problems balancing out the workloads over time. As things stand, it is common that, in order to achieve an acceptable quality-of-service at peak load, business servers on average run at only 5 to 10 percent load. Such an amount of unused resources translates into huge costs. These costs could be cut with the help of utility computing, where someone else employs the unused resources when they are not needed by the business.

What's more, having a fixed computing pool can become a straitjacket that stifles innovation of new services and damages the flexibility of a company.

For instance, we see more and more businesses that produce services on the Web. These new services are often the most innovative users of the new network and information technology that we see today. Network businesses are vulnerable to flash crowds – hundreds of thousands of users that access an online service at the same time – because it is uneconomical to dimension resources to deal with such intense load peaks. For such businesses to be reliable and viable, we need infrastructure that can allocate resources to these businesses in a short time, from a common resource pool.

Can a virtually infinite resource pool be built? Yes, but it will require an infrastructure that encourages investment in scarce resources. More investment in the right places enables the pool to grow organically, and to adapt to unknown workloads that will occur in the future and services that we cannot yet begin to imagine.

To begin creating an infrastructure that turns computing power into an almost liquid resource, SICS is building, in cooperation with HP Labs, the next-generation computing infrastructure 'Tycoon'. This provides a public, open market, where computing resources are sold for profit. The market is designed to encourage on-demand resource acquisition, rather than reservations. The motivation is that you rarely know your future demands in such detail that they could be easily expressed in a resource reservation language. Instead, you typically notice at run-time that more resources are needed to meet the quality-of-service objectives for response latency, throughput and so on.

The Tycoon market uses dynamic prices that fluctuate in response to changing supply and demand. Demand-driven prices make it more profitable to provide scarce resources, and encourage resources to be added to the network when a shortage occurs. The market is open, meaning resources are provided by competing providers; this prevents consumers from being locked in by a particular provider.

SICS is also active in the EU-funded SORMA project, led by the University of Karlsruhe. SORMA's goal is to create a truly open and 'self-organising' marketplace for the future computing grids. In contrast to Tycoon, which provides one specific computing platform built on virtualization, SORMA will be a general market platform or hub, that will connect together heterogeneous computing platforms in one common infrastructure. SORMA's members are active in many of the existing Grid computing initiatives worldwide, such as gLite, Mosix, and CATNETS. We expect these will benefit greatly from the market features that will be provided by the SORMA infrastructure.

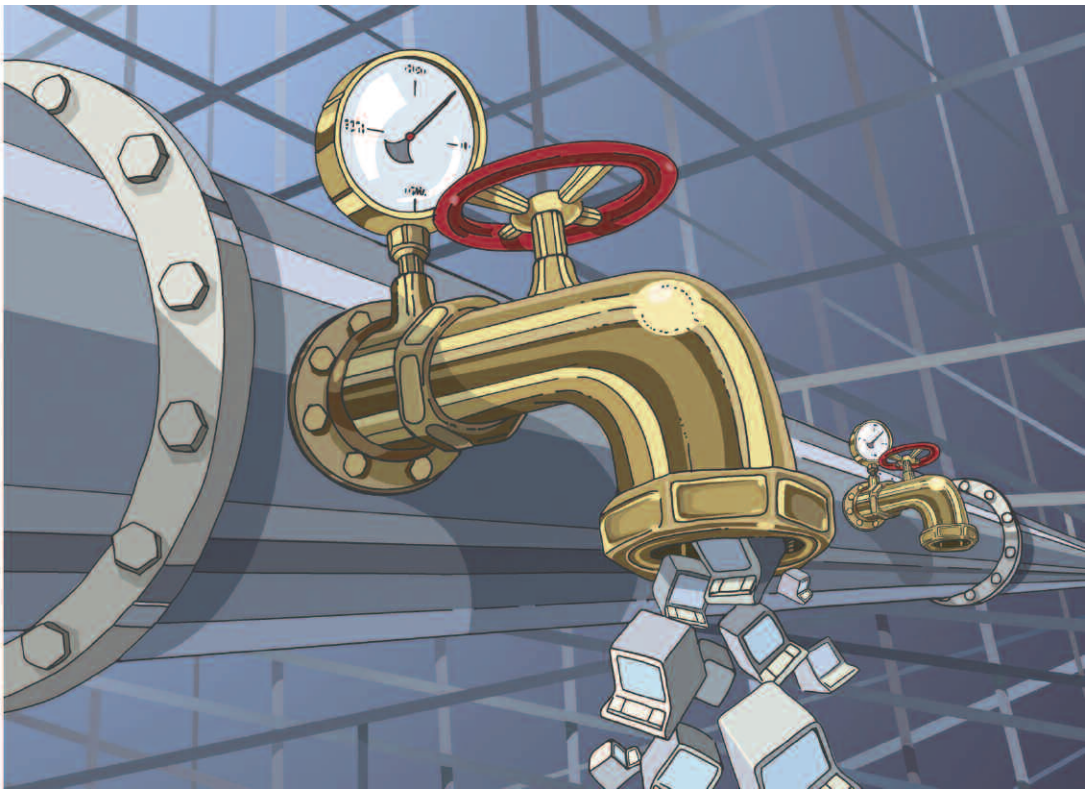

The next-generation computing infrastructure "Tycoon" is being built to turn computing power into an almost liquid resource.

Link:
http://tycoon.hpl.hp.com/

Please contact:
Lars.Rasmusson
SICS, Sweden
Tel: +46 8 633 15 00
E-mail: Lars.Rasmusson@sics.se

New Hungarian National Cancer Registry

by László Kovács and Máté Pataki

The aim of the National Cancer Registry GRID project (Országos Rákregiszter GRID - ORG) is to develop the next generation of the National Cancer Registry (NCR) for Hungary. The NCR started operation in 1999, and its central mission is the collection, management, and analysis of medical data on people who have been diagnosed with malignant or neoplastic disease (cancer). The technology behind NCR is now outdated, and our goal is to create a modern distributed system with a user-friendly Web-based interface and secure data transfer.

The ORG project is a consortium of the Department of Distributed Systems (DSD) of SZTAKI, Arvato Systems Hungary Inc. and the National Institute of Oncology. It has been responsible for building new online infrastructure to collect and validate medical data, which will greatly improve the quality of data on cancer diseases in the NCR and thus provide a much stronger statistical base for decision-makers and medical researchers.

One of the other important objectives of the ORG project is to broaden the range of data associated with cancer diseases. The current NCR collects only cancer-specific medical data and provides annual reports and statistics focused on the geographic spread of cancer cases. In the ORG system we plan to collect other healthcare, environmental, political, demographical and economical data associated with a given geographical territory (grid cell). Associating these data with the currently collected cancer-specific records will give a much more comprehensive view of possible causes of cancer, and aid experts in the analysis of factors affecting the progress of cancer diseases.

The quality of the collected data is an essential factor, not only for researchers but also for physicians. For instance, the latter could use the system to identify similar cancer cases, and thus for decision support. To achieve reliability in the collected data we introduced so-called consistency criteria. These are like small plug-ins, with each analysing incoming records for different types of errors, such as invalid address, name and sex mismatch, miscoding of disease code (BNO) etc. During the lifetime of the system new requirements could emerge or frequent errors in the incoming data be discovered. By writing consistency criteria and adding them to the system, further occurrences of these errors can be easily detected and avoided. These criteria will then warn the operators if an error is found or suspected. The operator or the physician can then decide if it is indeed an error, or merely an unusual case.

To make this system more accessible a Web-based interface is used, allowing the new registry to be accessed easily from any browser without having to install third-party software. Because data collected by the registry is sensitive and confidential, privacy and security are important issues. Access is therefore limited to the internal network of the hospitals, and even there the system uses a secure transfer protocol. Each hospital can install client software, which will automatically import the necessary data from their internal system (warning immediately if errors are found by either the local consistency criteria or the central server). Within this distributed system, these clients communicate with the server also by secure connections. This in itself is a huge step forward, since up until now, data was often sent on a floppy disk by regular post.

The ORG Cancer Registry is currently in a test phase, which will finish in early 2007. After this, the old and the new systems will run in parallel in order to eliminate any remaining errors in the system, train personnel, and prepare for the final switch from the old system to the new one.

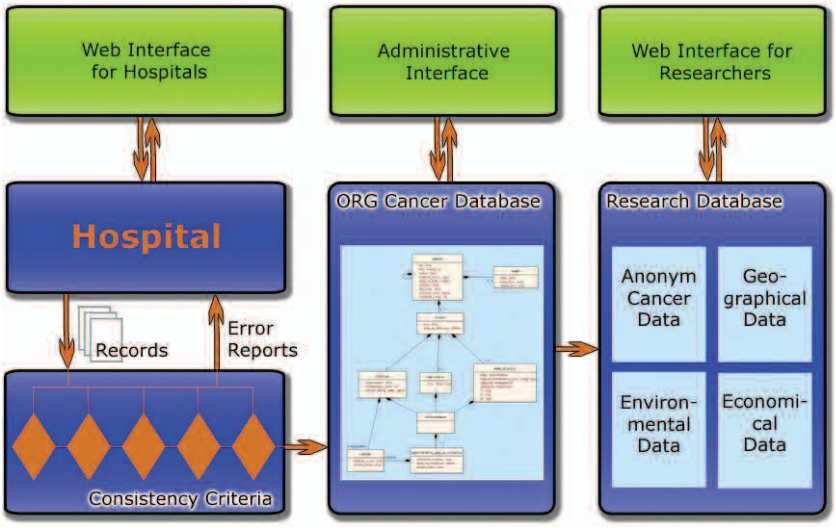

Architecture of the National Cancer Registry GRID.

Link:
ORG Project Page:
http://dsd.sztaki.hu/projects/org/en/

Please contact:
László Kovács
SZTAKI, Hungary
Tel: +36 1 279 6212
E-mail: laszlo.kovacs@sztaki.hu

Efficiency and Quality of Service of Internet Service Providers

by Oliver Heckmann

The Internet is a large network formed from 30,000 autonomous systems (AS), operated by thousands of Internet service providers (ISPs). While these ISPs compete with each other for customers and traffic, they must also cooperate and exchange traffic in order to maintain worldwide connectivity. In contrast to the traditional telecommunication markets, there are almost no central organisations in the Internet that enforce cooperation and regulate the market.

Given the competition among ISPs, it is vital that they operate their networks efficiently. Quality of service (QoS) is also important: many emerging multimedia applications such as voice and video communication can greatly benefit from QoS support in a network. This therefore opens further possibilities for value-added services with which providers can differentiate themselves and target new markets.

Hence, efficiency and QoS are crucial for ISPs operating the networks that comprise the Internet. Recent research, represented in the dissertation of Oliver Heckmann, gives new insights into how ISPs can improve in these areas. For his work, Oliver Heckmann was recently presented with the ERCIM Cor Baayen award. The central question of the dissertation is: how can an Internet service provider optimize the efficiency and quality of service of its network?

This work shows that efficiency and QoS strongly depend on three factors:

- the network architecture (how the network is built)
- traffic engineering and network engineering (how the network is adapted to changes)
- the interconnections (how the network is connected with other networks).

A system-oriented approach is therefore employed, in which all these areas are analysed and optimized, while taking into consideration their mutual influences (see figure).

In the context of network architectures, the focus of the research lies on methods for providing QoS. Analytical models are developed in order to determine and analyse the overprovisioning factor of a plain best-effort network compared to networks that support explicit service differentiation. The

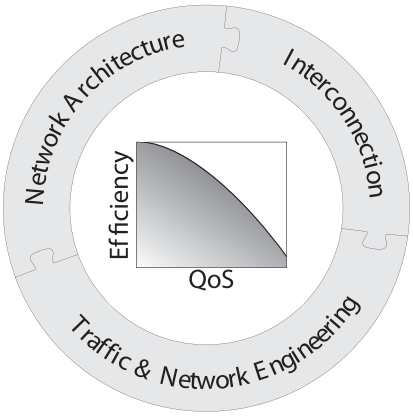

Efficiency and quality of service strongly depend on three factors.

work shows that the analytical overprovisioning factor ranges between three and five in most cases. Furthermore, the different IETF QoS architectures are evaluated and compared in an experimental study to confirm and extend previous findings.

In this context, a bandwidth broker for Differentiated Services networks is developed. It offers strong QoS guarantees and a high efficiency by overbooking. With the bandwidth broker, the so-called Charny bound can be broken and the utilization of networks with services of the highest QoS can be increased by a factor of about four.

The work shows that besides the network architecture, the interconnection structure also strongly influences the efficiency and QoS of an ISP. The interconnection structure describes how the network is connected with its peering and transit networks. Different approaches to optimizing the efficiency, reliability and QoS of the structure are described as decision problems, and are then solved and evaluated. The dissertation shows that with these approaches, cost savings of 5% to 30% can be realized, and QoS improved.
The effectiveness of using traffic engineering to further improve the efficiency and QoS obtained with the network architecture and interconnections is also analysed. Weaknesses in existing approaches are identified and corrected, taking the findings in the rest of the dissertation into account.

Due to the steadily increasing traffic volumes (Internet traffic is doubling in volume every 9-15 months), capacity expansion is the most important network engineering task of a provider. New strategies for capacity expansion are therefore elaborated and evaluated in this dissertation. They have significant advantages over the rules of thumb used today, since they also consider the influence of different network architectures and traffic engineering. The results show that traffic engineering should be explicitly taken into account when planning capacity expansions.

Links:
Online publication of the dissertation:
http://elib.tu-darmstadt.de/diss/000522/

Homepage of the book 'The Competitive Internet Service Provider' by Oliver Heckmann (contains selected parts of the dissertation, introductory chapters plus updates):
http://www.competitive-isp.info

Please contact:
Oliver Heckmann
E-mail: oliver.heckmann@gmail.com

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

CONCUR 2006 - The 17th International Conference on Concurrency Theory

by Christel Baier and Holger Hermanns

The 17th International Conference on Concurrency Theory was hosted by the University of Bonn, Germany, from 27-30 August 2006. There were 267 registered participants to CONCUR 2006 and its satellite events from nineteen countries around the world.

The purpose of the CONCUR conferences is to bring together researchers, developers and students in order to advance the theory of concurrency and promote its applications. Interest in this topic is growing continuously as a consequence of the importance and ubiquity of concurrent systems and their applications, and of the scientific relevance of their foundations. The scope covers all areas of semantics, logic and verification techniques for concurrent systems. Topics include concurrency-related aspects of models of computation and semantic domains, process algebras, Petri nets, event structures, real-time systems, hybrid systems, probabilistic systems, model checking, verification techniques, refinement techniques, term and graph rewriting, distributed programming, logic constraint programming, object-oriented programming, typing systems and algorithms, security, case studies, tools and environments for programming and verification.

The first two CONCUR conferences were held in Amsterdam (NL) in 1990 and 1991, with the locations in following years being Stony Brook (US), Hildesheim (DE), Uppsala (SE), Philadelphia (US), Pisa (IT), Warsaw (PL), Nice (FR), Eindhoven (NL), University Park (Pennsylvania, US), Aalborg (DK), Brno (CZ), Marseille (FR), London (UK) and San Francisco (US). The proceedings of CONCUR 2006 have appeared in Springer LNCS as volume 4137.

The technical programme of the conference and its satellite events delivered 122 presentations in total, from which fifteen were invited talks or invited tutorials. Out of 101 regular papers submitted to CONCUR 2006, 29 were accepted for presentation at the conference. The conference also included three talks by excellent invited speakers: Edward Lee (University of California at Berkeley, US), Orna Kupferman (Hebrew University, IL) and Jan-Willem Klop (Free University of Amsterdam, NL). Further, the programme included two plenary invited tutorials, delivered by Roberto Segala (University of Verona, IT) and Uwe Nestmann (Technical University of Berlin, DE).

The talk by Edward A. Lee on 'Making Concurrency', which was jointly invited by CONCUR and its satellite workshop 'Formal Methods for Industrial Critical Systems' (FMICS), was a particular highlight. FMICS is also the annual meeting of the ERCIM Working Group with the same title and was organised by Martin Leucker (TU Munich, DE) and Lubos Brim (Masaryk University, CZ). See ERCIM News 67 (http://ercim-news.ercim.org/content/view/24/46/) for a report on FMICS 2006.

Apart from FMICS, CONCUR 2006 was accompanied by the following spectrum of satellite workshops:

- Verification of Infinite-State Systems (INFINITY), organised by Ahmed Bouajjani (University of Paris 7, FR)
- Expressiveness in Concurrency (EXPRESS), organised by Iain Phillips (Imperial College, UK) and Roberto Amadio (University of Paris 7, FR)
- Structural Operational Semantics (SOS), organised by Rob van Glabbeek (National ICT Australia, AU) and Peter Mosses (Swansea University, UK)
- Geometric and Topological Methods in Concurrency (GETCO), organised by Eric Goubault (CEA, FR)
- German Verification Day (GVD), organised by Juergen Niehaus (University of Oldenburg, DE)
- Foundations of Coordination Languages and Software Architectures (FOCLASA), organised by Carlos Canal (University of Malaga, ES) and Mirko Viroli (University of Bologna, IT)
- Parallel and Distributed Model Checking (PDMC), organised by Boudewijn Haverkort (University of Twente, NL) and Jaco van de Pol (CWI, NL)
- Control and Observation of Real-Time Open Systems (CORTOS), organised by Franck Cassez (CNRS, FR)
- Graph Transformation for Verification and Concurrency (GT-VC), organised by Arend Rensink (University of Twente, NL), Reiko Heckel (University of Leicester, UK) and Barbara Koenig (University of Stuttgart, DE).

The conference took place in Germany for the second time after 1993. The organisers attempted to give the conference additional stimuli through moderate conference fees, making it attractive to participants from Central and Eastern Europe and to young researchers all over the world.

CONCUR 2006 was sponsored by ERCIM.

Link:
http://d.cs.uni-sb.de/concur2006

Please contact:
Holger Hermanns
Saarland University, Germany
E-mail: hermanns@cs.uni-sb.de

25 Years of the IFSR: A Milestone in the Systems Sciences

by Gerhard Chroust

The terrible events between 1914 and 1945, World War I, the world economic crisis and World War II, induced scientists like Ludwig von Bertalanffy, Norbert Wiener, W. Ross Ashby and their colleagues to offer a way out: holistic rather than fragmented thinking. This led to two new sciences: systems theory and cybernetics. In 1980 the International Federation for Systems Research (IFSR) was founded as a federation of all systems societies. The celebrations of its 25th anniversary provided a chance to look at the successful history of the IFSR and at the current and future state of the systems sciences.

The celebration of the 25th anniversary of the IFSR were held in Vienna on 19 April 2006 during the European Meeting of Cybernetics and Systems Research. The opening presentation was given by Secretary General Prof. Gerhard Chroust (University of Linz, Austria), who described the aims, goals, structure, projects and achievements of the IFSR. The designated president of the IFSR, Prof. em. Matjaz Mulej, University of Maribor, Slovenia, then gave the traditional Ross Ashby Memorial Lecture. The rest of the day was devoted to a review of past and present activities of the IFSR and an outlook into the future.

Matjaz Mulej chose the title 'Systems theory – a worldview and/or a methodology', and argued that the immensely growing knowledge of humankind causes unavoidably narrow specialization of individuals, with every profession working in a partial field. Knowing and using parts alone rather than wholes is not enough. Warnings about the problems of oversights, one-sidedness and their consequences were voiced millennia ago, but have been neglected by specialists in a variety of professions who fail to recognize that specialization is not sufficient.

As Ludwig von Bertalanffy (1901-1972) – the father of general systems theory – emphasized, systems thinking is not a profession but a world view. It is a matter of education, values, culture, ethics, and norms of behaviour, added to every single field of professional knowledge.

The IFSR was based on the recognition that in order to respond to global problems, synergistic cooperation between all scientific societies working on systems theory and cybernetics is necessary. With support from the Austrian Federal Ministry for Science and Research, it was founded on 12 March 1980, with three members. These were the Society for General Systems Research, now ISSS (its president, Prof. J. Klir of the USA, became the first president of the IFSR); the Austrian Society for Cybernetic Studies (its president, Prof. Robert Trappl of Austria, became the first vice-president); and the Systeemgroep Nederland (its representative, Prof. Gerard de Zeeuw of Netherlands, became the first secretary treasurer).

Left to right: Gerhard Chroust, secretary general; Matjaz Mulej, president; Gerard de Zeeuw past secretary/treasurer and past president; and Gary Metcalf, vice-president.

The aims of the federation were to stimulate activity associated with the scientific study of systems, and to coordinate such activities on the international level. Since 1980 the federation has grown to 32 members from 25 countries on all continents. Past and present major contributions to the field of systems research are:

- the Journal of Systems Research and Behavioral Science (Editor-in-chief: Prof. M. C. Jackson), publishing original articles on new theories, experimental research and applications relating to all levels of living and non-living systems
- the IFSR International Book Series on Systems Science and Engineering (editor-in-chief: Prof. George Klir), founded in 1985, publishing high-quality scientific papers on systems sciences
- the Fuschl Conversations (started in 1980 by Bela H. Banathy), which provide an alternative to conventional conferences by offering a collectively guided disciplined inquiry into issues of social/societal significance by scholarly practitioners in self-organised face-to-face teams. Since 1980 they have been held biannually in Fuschl, near Salzburg, Austria and replicated in many other places to an estimated total of forty or fifty
- the First International Congress of the IFSR (14-17 November 2005 in Kobe, Japan) under the chaimanship of Prof. Jifa Gu, Prof. Y. Nakamori, and Prof. G. Chroust, was the first international systems conference to be held in Japan. It established a foundation for understanding and collaboration between Japanese and non-Japanese systems professionals
- the Ross Ashby Memorial Lecture at the bi-annual EMCSR-Conferences, sponsored by the IFSR
- a Web site and newsletter (Editor-in-chief and Webmaster: Prof. Magdalena Kalaidjieva, recently succeeded by Prof. G. Chroust).

Recent and current projects of the IFSR include:

- acquiring, archiving and scientifically analysing the Ludwig von Bertalanffy Legacy in cooperation with the Bertalanffy Centre for the Study of Systems Science, Austria (Prof. Wolfgang Hofkirchner, Salzburg)
- an electronically accessible International Encyclopaedia of Systems and Cybernetics based on Charles François' book version (Prof. Günther Ossimitz)
- planning an Academy of Systems Sciences and Cybernetics (Prof. Jifa Gu and Prof. Matjaz Mulej)

- collecting and analysing data on courses in systems sciences to develop a basic common core (Prof. G. A. Swanson)
- networking the archives of pioneers of systems science and cybernetics (Prof. Gerhard Chroust, Gary Metcalf);
- connecting systems scientists in Asia and the West (Prof. Jifa Gu).

The anniversary event was recorded and the video will be published on DVD in 2007.

Link:
http://www.ifsr.org

Please contact:
Gerhard Chroust
Kepler University, Linz, Austria,
E-mail: gc@sea.uni-linz.ac.at

Matjaz Mulej
Univerza v Mariboru, Slovenia
E-mail: mulej@uni-mb.si

PRACTICE AND THEORY OF AUTOMATED TIMETABLING

PATAT 2006 Conference

by Hana Rudová

The 6th International Conference on Practice and Theory of Automated Timetabling (PATAT) was organised by the Faculty of Informatics at Masaryk University, a member of the Czech Research Consortium for Informatics and Mathematics (CRCIM). The conference took place in Brno, Czech Republic from 30 August to 1 September and was attended by almost a hundred delegates from 23 countries.

The PATAT conferences, which are held every two years, bring together researchers and practitioners in all aspects of computer-aided timetable generation. This includes university timetabling, school timetabling, personnel rostering, transportation timetabling, sports scheduling and other areas of the subject. These timetabling problems are common to many of us. The classical example is the construction of the timetable for some school. It corresponds to an assignment of the time and room to each course. Such a generated timetable must satisfy requests given by student programs, teachers and students. Another important problem can be introduced by timetabling of the nurses at hospitals. Here the goal is to find nurses for each shift such that the needs of the hospital and its employees are satisfied. A very commom problem, solutions for which were also explored by the conference, is sport league or tournament scheduling. The goal is to find team-mates, dates and placement of the games within constraints such as the availability of venues, the order of opponents, or travel time and distance.

PATAT conference participants.

The conference in the Czech Republic was a very successful continuation of former conferences organised in the USA, Belgium, Germany, Canada and the UK. Programme chairs of the conference were Edmund Burke from the University of Nottingham and Hana Rudová from the organising university. The programme of this year's conference featured seventy presentations representing the state of the art in automated timetabling. This included four plenary papers, seventeen full papers, 41 extended abstracts, and eight system demonstrations. Plenary presentations discussed physician scheduling in emergency rooms (Michel Gendreau, Université de Montréal), differences between theory and practice in university timetabling (Barry McCollum, Queen's University and eventMAP Ltd), search techniques in timetabling (James Orlin, MIT), and state of the art in timetabling experimentation (Andrea Schaerf, Università di Udine).

The conference proceedings were published by Masaryk University Press and selected and revised papers will be published in Lecture Notes in Computer Science by Springer-Verlag.

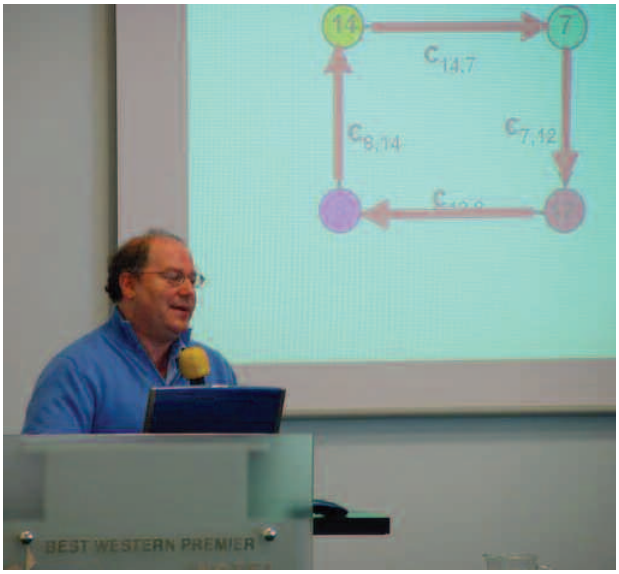

Invited speaker James Orlin from the MIT talking about Very Large-Scale Neighborhood Search Techniques in Timetabling Problems.

Link:
http://patat06.muni.cz

Please contact:
Hana Rudová
Faculty of Informatics, Masaryk University
E-mail: hanka@fi.muni.cz

ESFORS – NESSI Workshop Addresses New Research Priorities in Trust, Security and Dependability in Software Systems and Services

by Aljosa Pasic and Jim Clarke

The emerging service centric ICT paradigm is changing the way infrastructure and applications will be managed and delivered in the next few years. The new challenges brought about by this paradigm and, in particular, the underlying security considerations are the main focus of the IST FP6 coordination action ESFORS (European Security Forum for Web services, software and systems), as well as of the working group "Trust, Security and Dependability" of the NESSI (Networked European Software and Service Initiative) European technology platform. The recent workshop held in Paris on 6-8 September 2006 shows that there is growing mutual interest on research in security and software engineering. The workshop was attended by more than 70 participants coming both from industry and academia. This article presents the main workshop conclusions as well as related technology roadmapping activities.

ESFORS, a Coordination Action that brings together the right stakeholders

ESFORS is a coordination action of the Information Society Technologies (IST) programme of the 6th framework programme (FP). It aims at bringing together the European stakeholders for secure and dependable Information and Communication Technologies (ICTs) to address the security and dependability requirements of emerging software systems and service platforms, such as web services.

The project will complement already existing coordination actions in the area of ICT for security, dependability and trust, such as SecurIST (http://www.securitytaskforce.eu/). It will co-operate with SecurIST to ensure that open service requirements are incorporated into the SecurIST security and dependability technology roadmap and that the roadmap is incorporated into the research agenda of the software and service research community, notably through the strategic research agenda of the NESSI technology platform. Moreover, ESFORS will act as a bridge between two research communities: the software and service engineering community and the security and dependability community.

In order to take into account the views and experience of the experts and stakeholders in the field of security and dependability of software systems and services, ESFORS has created an expert group of well known researchers and experts from academia, industry and research institutes from nearly all states of the European Union (EU).

ESFORS organises conferences and workshops that are the main vehicle for giving an opportunity to the research communities and expert groups to present their ideas, discuss and reach consensus upon the future research challenges in security and dependability that need to be addressed in prospective European research programmes. The workshops are focused on particular topics that deal with security and service oriented software architectures and infrastructures.

European Technology Platform NESSI and NESSI Working Groups

Promoted by thirteen major European ICT corporations, NESSI aims to develop a visionary strategy for software and services driven by a common European research agenda. NESSI has been officially launched on 7 September 2005 in Brussels. During the last NESSI general assembly, held on 8 June 2006 in Brussels, there has been an opportunity to present the working groups as well as the NESSI contribution to the research priorities of the EU's 7th FP. New partners and new community members have also been announced during that event. The core partners of ESFORS proposed and established one of the first NESSI working groups to officially receive approval, entitled 'Trust, Security and Dependability Working Group' (TSD WG).

Research Priorities for Trust, Security and Dependability in Software Systems and Services

Software applications will be broken into separately managed component services and will form so called dynamic service coalitions. This has many consequences from a security point of view: first, applications will need to utilise components out of different domains of control that are required to obey separate security policies and ask for diverse security and dependability qualities; second, components may be owned and operated by different organisations so that informal company arrangements will have to be replaced by formal agreements resulting from partially automated negotiations; and third, the services will be shared between many consumers, which implies advanced confidentiality, virtualization and isolation requirements.

Taking these evolutions into consideration, the participants of the ESFORS-NESSI workshops held in Paris on 6-8 September 2006 have derived the following conclusions regarding the main research priorities to focus on (this is a non exhaustive list):

- *Privacy considerations* - a clear definition of privacy as well as research scope is needed in order to unify criteria and to decide on the most appropriate research directions. Some of the suggested topics to further explore and develop include: User empowerment, usage control and visibility, context-aware privacy, data issues in software services, protection control and accountability.
- *Identity considerations* - after discussion on the identity concept, the existence of a unique or multiple identities per individual in service-oriented environments, the possibility to manage identity in dynamic service coalitions and many other topics have been debated. These include: links between privacy, reputation, identity, credential systems and interoperability, solutions for 'special' identity cases (group identity, inclusive identity, identity of “things” etc), virtualization of identity, interoperability, identity control lifecycle (issue, verification, revocation, etc.), repudiation of identity etc.

- *Security mechanisms for services* - this is an extensive research topic since it encompasses issues like: clear division between modification of old security mechanisms and establishing new mechanisms, policies integration and enforcement, negotiation and conflict resolution mechanisms, adaptive security, reusability and portability of mechanisms, etc.
- *Trust analysis, management and monitoring* - these are two different, but inter-related, sub-topics that include, respectively: Trust analysis and management including trust modelling and computational trust; and dependability assessment and monitoring.
- *Security and dependability engineering* - some issues concerning security and dependability engineering are: use of patterns, integration of 'lessons learnt' and formal models, semantics, secure coding, secure migration, etc. There was a particular emphasis on decision making in security engineering (user-driven, risk-driven, threat-driven, trust-driven, complexity- and cost-driven, etc).
- *Stakeholders expectations* - a number of key issues were discussed from the perspective of various stakeholders including end users, enterprises, academia, and service and communications providers. These topics included challenges related to the risk management approach to security, including quantitative risk assessment, the need for methods and tools for complex scenarios, raising awareness, the relationship between trust and training, and costs vs. benefits across all three perspectives: citizen, industry and regulatory. In addition, the need for secure metrics and monitoring protocols and taxonomy or common language are needed to ensure that all stakeholders on security can communicate in a mutually understand language, which is very difficult in this environment.

The full report presenting in detail the conclusions of the ESFORS-NESSI workshop held in Paris can be found at http://www.esfors.org/index.php?option=com_remository&Itemid=65.

The above research topics and priorities were further examined with research stakeholders in the NESSI TSD WG in two meetings held in September and November 2006, respectively. In addition, the ESFORS project organised and hosted a networking session at the IST 2006 event on 23 November 2006 in Helsinki. The well-attended session was entitled 'Bringing Together the European Stakeholders for Secure Service and Software Engineering: Creating trust, security and dependability in service applications'. The keynote addresses for the ESFORS networking session and the project proposal presentations can be found at the ESFORS web site. The next ESFORS-NESSI workshop is scheduled for June 2007 in Maribor, Slovenia.

Links:
ESFORS: http://www.esfors.org
NESSI: http://www.nessi-europe.com/

Please contact:
ESFORS coordinator Pedro Soria, Atos Origin
E-mail: pedro.soria@atosorigin.com

NESSI WG chairman Aljosa Pasic, Atos Origin
E-mail: aljosa.pasic@atosorigin.com

CALL FOR PAPERS

FCT 2007 - 16th International Symposium on Fundamentals of Computation Theory

Budapest, Hungary, 27-30 August 2007

The Symposium on Fundamentals of Computation Theory was established in 1977 for researchers interested in all aspects of theoretical computer science, in particular in algorithms, complexity, formal and logical methods. It is a biennial series of conferences previously held in Poznan (1977), Wendisch-Rietz (1979), Szeged (1981), Borgholm (1983), Cottbus (1985), Kazan (1987), Szeged (1989), Gosen-Berlin (1991), Szeged (1993), Dresden (1995), Kraków (1997), Iasi (1999), Riga (2001), Malmö (2003), and Lübeck (2005).

Authors are invited to submit papers presenting original unpublished research in all areas of theoretical computer science. Topics of interest include (but not limited to):
- automata and formal languages
- design and analysis of algorithms
- computational and structural complexity
- semantics
- logic, algebra and categories in computer science
- circuits and networks
- learning theory
- specification and verification
- parallel and distributed systems
- concurrency theory
- cryptography and cryptographic protocols
- approximation and randomized algorithms
- computational geometry
- quantum computation and information
- bio-inspired computation.

Invited speakers:
- Ahmed Bouajjani (Paris, France)
- Oscar H. Ibarra (Santa Barbara, CA, USA)
- László Lovász (Budapest, Hungary)
- Philip Scott (Ottawa, Canada).

Important dates:
- Deadline for submissions: 5 March 2007
- Notification to the authors: 20 April 2007
- Final version: 20 May 2007.

The proceedings will be published in the Lecture Notes in Computer Science series of Springer-Verlag and it will be distributed at the conference. We anticipate that a special issue of Theoretical Computer Science will be devoted to selected papers published at the conference.

More information:
http://www.conferences.hu/fct2007/

CALL FOR PAPERS

PACT 2007 - 16th International Conference on Parallel Architectures and Compilation Techniques

Brasov, Romania 15-19 September 2007

The purpose of PACT is to bring together researchers from architecture, compilers, applications and languages to present and discuss innovative research of common interest. PACT features cutting-edge research on a broad range of topics, that include, but are not limited to:

- parallel architectures and computational models
- compilers and tools for parallel computer systems
- multicore, multithreaded, superscalar, and VLIW architectures
- compiler/hardware support for hiding memory latencies
- support for correctness in hardware and software (esp. with concurrency)
- reconfigurable computing
- dynamic translation and optimisation
- I/O issues in parallel computing and their relation to applications
- parallel programming languages, algorithms and applications
- middleware and run-time system support for parallel computing
- high performance application specific systems
- applications and experimental systems studies
- non-traditional computing systems topics.

Important dates:

- 26 March 26, 2007: abstract submission deadline
- 1 April 2007: paper submission deadline
- 2 April 2007: workshop proposal submission deadline.

More information:
http://www.pactconf.org/pact07

Fellowships available in GRID Research

The CoreGRID Network of Excellence currently offers fellowships for postgraduate students in the field of GRID Research The CoreGRID web site also offers the possibility to post job announcements related to GRID research. Job postings are free of charge for academic institutions and organisations. CoreGRID is a Network of Excellence administrated by ERCIM.

More information;
http://www.coregrid.net/jobs

Advertisement

In 2007, INRIA will recruit 600 new staff

Take up the challenge of working with the very latest technologies in your profession:

research, computing, finance, assistantship, communication, HR, industrial relations, logistics...

Find out more about research subjects, activity schedules, the list of available positions, and download the application files:

It starts here and now...
www.inria.fr/travailler

Researchers: 125 posts
Candidates will be recruited for secondment or delegation positions through a competitive selection process with an aim to strengthen INRIA's research capacity within the scope of its scientific challenges:
- network
- software security
- complex systems
- simulation and virtual reality modelling living systems

The **competitive selection process** is open to both young and experienced researchers, in France and abroad.

Start of the recruitment campaign:
15 December 2006

R&D researchers: 150 posts with fixed-term contracts
These 1-3 year contracts are designed for **young graduates** (at least five year's higher education), **specialists** from industry and **foreign specialists** from the academic world.
They can be used to acquire professional experience or to bring expertise INRIA research teams and to advanced research support services.
Recruitment: research contracts awarded throughout the entire year

Young graduate recruitment campaign:
15 March 2007.

Support staff: 90 posts
Covering the entire professional spectrum: legal, management, finance, accounting, HR, assistantships, documentation, communication, logistics, etc.
These posts are aimed at young graduates and experienced professionals.

Start of the recruitment campaign:
Civil servant transfers: 2 January 2007
External competitive selection process: end of March 2007

Not forgetting
Training-through-research: 460 PhD students and post-docs
460 positions are available in research teams. Experience abroad - post-docs will be in particularly high demand.
Recruitment: throughout the entire year with 2 busy periods
PhD students: 1 March 2007
Post-docs: 1 February 2007

INRIA is the French National Institute for Research in Computer Science and Control. It is a major European player specialised research institute in IT Research. 3,600 people - including 2,800 scientists (INRIA and partners) - work in its research units that are located in 7 regions in France.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE | INRIA

SPONSORED BY ERCIM

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

Trondheim, Norway 11-15 June 2007

'Ubiquitous Information Systems Engineering' is the theme of the conference this year. Most modern information systems span activities performed in several organisations and at different geographical locations. They often support the untethered mobility of their users. The systems have already today a large impact on every day life in the organisations, and on the individuals. As we move towards ambient, pervasive and ubiquitous computing this impact will increase significantly.

CAiSE'07 aims to bring together practitioners and researchers in the field of information systems engineering. CAiSE'07 invites submissions on the development, maintenance, procurement and usage of information systems - especially submissions dealing with aspects related to information systems engineering in ubiquitous environments.

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

SPONSORED BY ERCIM

ECOOP '07 - 21st European Conference on Object-Oriented Programming

Berlin, 30 July - 3 August 2007

The ECOOP 2007 conference will feature high quality papers presenting research results or experience in all areas relevant to object technology, including work that takes inspiration from or builds connections to areas not commonly considered object-oriented. Many different research methods can be applied, eg, both experimentally based work and mathematical results.

ECOOP 2007 will also host a number of workshops addressing different areas of object-oriented technology. Workshops serve as a forum for exchanging late breaking ideas and theories in an evolutionary stage. They typically focus on either in depth analysis or broad-ranging approaches to areas related to object-oriented technology. The conference also is offering a number of opportunities for tutorials and demonstrations of research and production systems.

More information:
http://ecoop07.swt.cs.tu-berlin.de/index.html

Editorial Information

ERCIM News is the magazine of ERCIM. This issue has a circulation of 10,500 copies. The printed version of ERCIM News has a production cost of €8 per copy. It is currently available free of charge.

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

Editorial Board:
Central editor:
Peter Kunz, ERCIM office (peter.kunz@ercim.org)
Local Editors:
Austria: Erwin Schoitsch, (erwin.schoitsch@arcs.ac.at)
Belgium:Benoît Michel (michel@tele.ucl.ac.be)
Czech Republic:Michal Haindl (haindl@utia.cas.cz)
Finland: Pia-Maria Linden-Linna (pia-maria.linden-linna@vtt.fi)
France: Bernard Hidoine (bernard.hidoine@inria.fr)
Germany: Michael Krapp (michael.krapp@scai.fraunhofer.de)
Greece: Eleni Orphanoudakis (eleni@ics.forth.gr)
Hungary: Erzsébet Csuhaj-Varjú (csuhaj@sztaki.hu)
Ireland: Ray Walsh (ray@computing.dcu.ie)
Italy: Carol Peters (carol.peters@isti.cnr.it)
Luxembourg: Patrik Hitzelberger (hitzelbe@lippmann.lu)
Norway: Truls Gjestland (truls.gjestland@ime.ntnu.no)
Spain: Salvador Lucas (slucas@dsic.upv.es)
Sweden: Kersti Hedman (kersti@sics.se)
Switzerland: Harry Rudin (hrudin@smile.ch)
The Netherlands: Annette Kik (Annette.Kik@cwi.nl)
United Kingdom: Martin Prime (M.J.Prime@rl.ac.uk)
W3C: Marie-Claire Forgue (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 office@ercim.org

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

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

A Joint INRIA-Microsoft Research Laboratory

INRIA and Microsoft have inaugurated a joint fundamental research laboratory in Orsay near Paris on 11 January 2007. For several years, INRIA and Microsoft Research have been working together to push forward research in the software industry field, a key sector in the 'weightless economy' that was the subject of a recent report submitted to the French ministry of economy.

The joint laboratory will employ around 30 top-level researchers who will devote their time to two key fields in computer science research and, more broadly, scientific research: using formal methods to demonstrate long theorems, or to validate communication protocols, and creating a new avenue of research in computer science dedicated to the interaction of significant volumes of data generated by modern scientific experimentation. The results of this research will be made public and accessible to the scientific community.

Link:
http://www.msr-inria.inria.fr/index.html

Keith Jeffery re-elected ERCIM President

Keith Jeffery, Director, IT & International Strategy at CCLRC, has been re-elected ERCIM president for a further

Keith Jeffery at the ERCIM meetings in Nice.

two-year term of office 2007-2009 during the bi-annual meeting of ERCIM's board of directors in Nice, France, in October 2006. Keith took office as ERCIM president on 1 January 2005.

Link:
http://www.cis.rl.ac.uk/Person/NAME=K.G.Jeffery

Colloquium in Memory of Gilles Kahn

Gilles Kahn was the Chief Executive Officer of INRIA and INRIA's representative on ERCIM's board of directors when he passed away in February 2006.

His friends and colleagues met on 12 January 2007 in Paris to commemorate his friendship and his scientific leadership.

Speakers at the colloquium included Samson Abramsky (University of Oxford), Nicholas Ayache (INRIA),

© INRIA / Photo Jim Wallace

Gilles Kahn in 1999

Gérard Berry (Esterel Technologies), Luca Cardelli (Microsoft Research, Cambridge), Gérard Huet (INRIA), Paul Klint (CWI), Pierre-Louis Lions (College of France), David Macqueen (University of Chicago), Robin Milner (University of Cambridge and École polytechnique), Gordon Plotkin (University of Edinburgh), Erik Sandewall (Linköping University), Laurent Théry (INRIA), and Jean Vuillemin (INRIA and École Normale Supérieure).

The first computer scientist to be elected to the French Academy of Science, a relentless promoter of science and supporter of ERCIM, Gilles Kahn was involved in many partnerships with the most prestigious organisations in the world. He was a great man of science.

He was and will remain, for those that were close to him, a warm and charismatic man whose simplicity and human qualities will remain in their memories and in their hearts.

Link:
http://www.inria.fr/gilleskahn/colloquium.en.html

ERCIM signed Berlin Declaration on Open Access

ERCIM has signed the 'Berlin Declaration on Open Access to Knowledge in the Sciences and Humanities' on 30 October 2006. The Berlin Declaration states that, "Establishing open access as a worthwhile procedure ideally requires the active commitment of each and every individual producer of scientific knowledge and holder of cultural heritage".

Open access contributions include original scientific research results, raw and metadata as well as source materials and digital representations of pictorial and graphical materials. The authors and right holders grant to all users a free, irrevocable, and universal right of access to these contributions and allow their work to be used, reproduced, or disseminated in digital form (provided correct attribution of authorship or copyright owner is given). Together with supplemental materials and the declaration of the right of use, the complete version of the work is to be made accessible in at least one electronic online archive.

The Berlin Declaration is in accordance with the spirit of the 'Bethesda Declaration on Open Access Publishing' and the 'Budapest Open Access Initiative'. Both also endorse fundamental changes to the practice of scientific publishing.

The Berlin Declaration was preceded by a conference at the the Max Planck Society in Berlin in 2003 where leading, international experts discussed new ways of accessing scientific knowledge and cultural heritage via the Internet. By January 2007, 193 organisations from all over the world have signed the declaration including the individual ERCIM members CWI, Fraunhofer-Gesellschaft, INRIA,

Links:
Berlin Declaration:
http://oa.mpg.de/openaccess-berlin/berlindeclaration.html

ERCIM News 64, European Scene: The Routes to Open Access Publishing:
http://www.ercim.org/publication/Ercim_News/enw64/

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Swiss Association for Research in Information Technology
c/o Professor Daniel Thalmann, EPFL-VRlab,
CH-1015 Lausanne, Switzerland
Tel +41 21 693 5214, Fax +41 21 693 5328
http://www.sarit.ch/

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

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

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

Order Form

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

Send, fax or email this form to:

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

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

I wish to subscribe to the

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

Name:

Organisation/Company:

Address:

Postal Code:

City:

Country

E-mail:

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