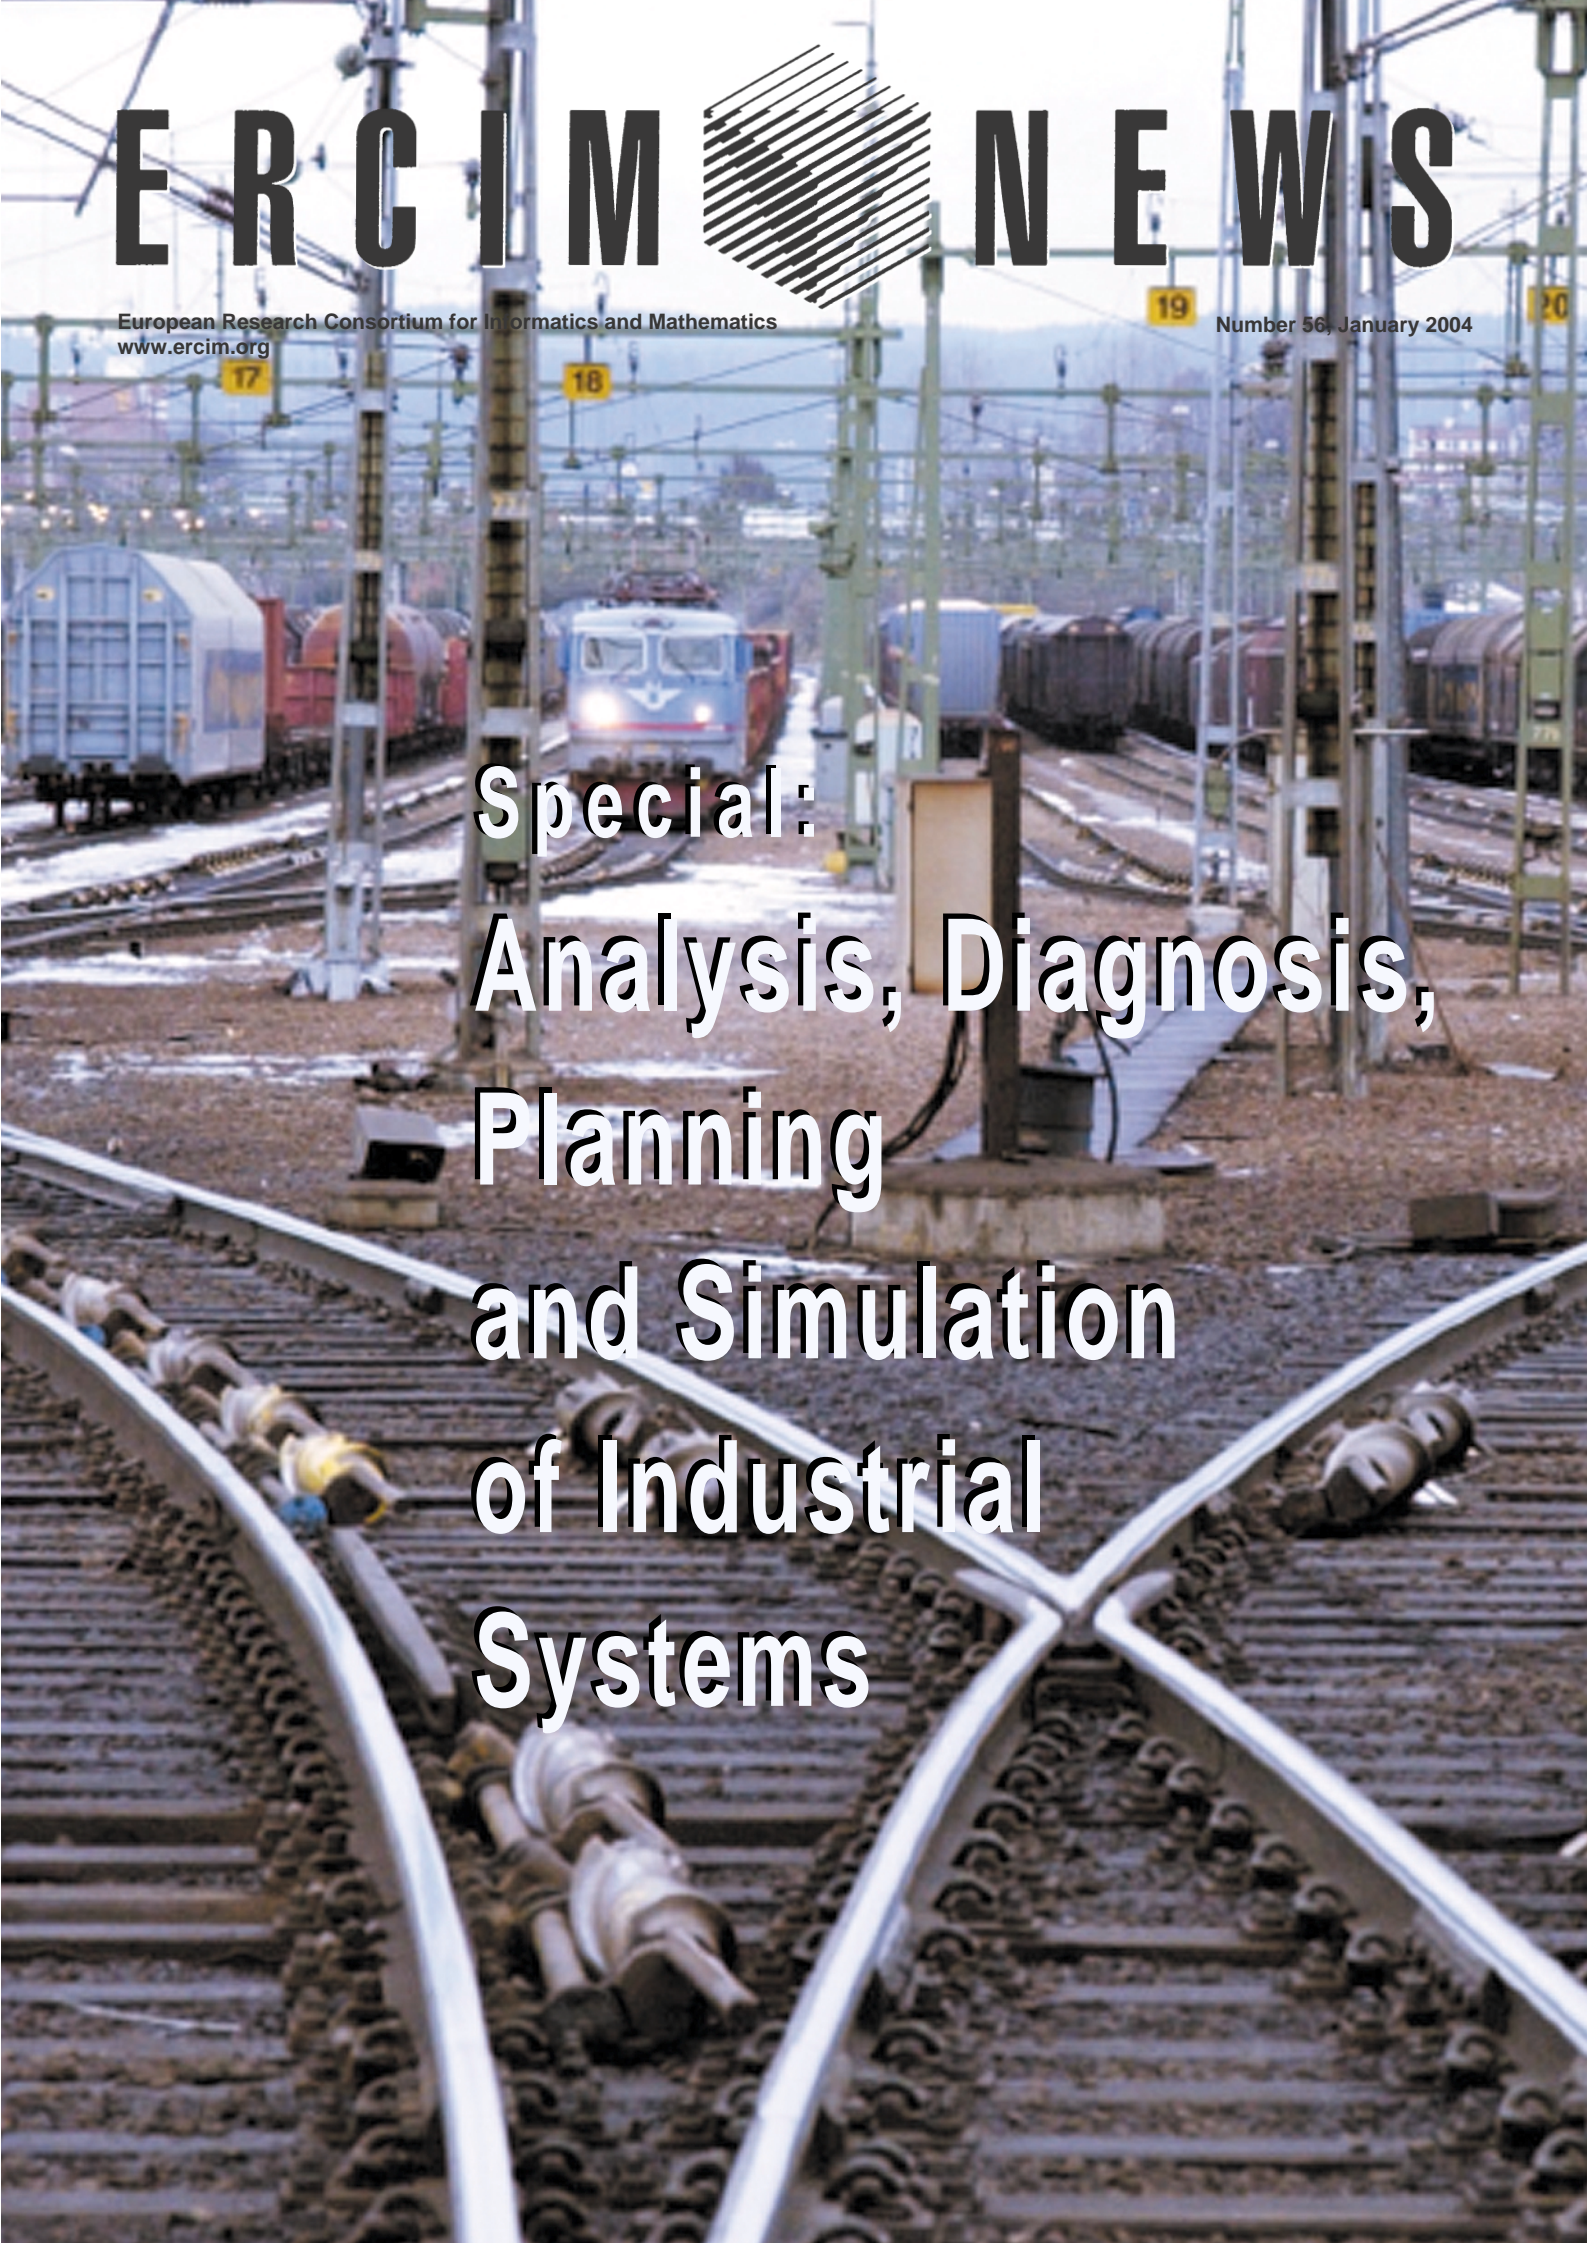


ERCIM NEWS

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

Number 56, January 2004



**Special:
Analysis, Diagnosis,
Planning
and Simulation
of Industrial
Systems**

Contents

JOINT ERCIM ACTIONS

- 3 **Keynote**
by Stelios Orphanoudakis, ERCIM President
- 4 **Tribute to Gerard van Oortmerssen**
by Stelios Orphanoudakis, ERCIM President
- 4 **Ronald de Wolf - Winner of the 2003 Cor Baayen Award**
- 5 **Possibilities and Limitations of Quantum Computing**
- 7 **First ERCIM Workshop on Informatics and Mathematics Applied to Interventional Medicine**
by Marc Thiriet, INRIA
- 8 **News from W3C**

SPECIAL THEME

- 10 **Industrial Diagnosis, Planning and Simulation — Introduction**
by Per Kreuger, SICS
- 12 **Efficient Algorithms for Inverse Magnetic Field Problems**
by Domenico Lahaye, CWI
- 13 **Deviation Detection of Industrial Processes**
by Anders Holst, Jan Ekman and Daniel Gillblad, SICS
- 14 **Bidding Agents Acquire Cargo in Online Auctions**
by Pieter Jan 't Hoen, CWI
- 15 **Decision Support System for Process Analysis and Supervision**
by Galia Weidl, University of Stuttgart, Germany
- 17 **Mathematical Challenges in the Electronics Industry**
by Wil Schilders, Eindhoven University of Technology, The Netherlands
- 18 **Contpack — Maximising the Volume Utilisation of Containers**
by Onno Garms, Ralf Heckmann and Stefan Rank, Fraunhofer SCAI
- 20 **First Commercial Bayesian Software for Intelligent Troubleshooting and Diagnostics**
by Claus Skaanning, Dezide, Denmark
- 22 **Innovating Shoe Manufacturing using Advanced Simulation Techniques**
by Emanuele Carpanzano and Andrea Cataldo, ITIA-CNR
- 23 **Billy Goat Detects Worms and Viruses**
by James Riordan and Diego Zamboni, IBM Zurich Research Lab/SARIT
- 24 **Decision Support for Personnel and Resources Management in Multi-Project Environments**
by Antonia Bertolino, Eda Marchetti and Raffaella Mirandola, ISTI-CNR
- 26 **Towards Useful Tools for Estimating Railway Infrastructure Capacity**
by Jan Ekman, SICS
- 27 **Efficient Utilisation of Resources in the Railway Industry by Levelling Departures and Arrivals of Trains**
by Martin Aronsson, SICS
- 28 **Real-Time Optimisation of Empty Rail Freight Car Distribution**
by Martin Joborn, Carmen Consulting, Sweden
- 29 **KUORMA: A Collection of APS-Algorithms for Forest Industry Wood Transport**
by Juha Savola, VTT
- 31 **Parc Bandwidth on Demand**
by Helmut Simonis, Parc Technologies Ltd, UK
- 32 **DesParO - A Design Parameter Optimisation Toolbox using an Iterative Kriging Algorithm**
by Christof Bäuerle, Clemens-August Thole and Ulrich Trottenberg, Fraunhofer SCAI

- 33 **On-line Testing of the Reactor Protection System in the Paks Nuclear Power Plant**
by István Varga, Tamás Bartha and Alexandros Soumelidis, SZTAKI
- 35 **Process Systems: Theory and Applications from Different Aspects**
by Katalin Hangos and Gábor Szederkényi, SZTAKI
- 37 **Multi-Agent System Technology in a Port Container Terminal Automation**
by Vicent J. Bottí, Universidad Politécnica de Valencia/SpaRCIM
- 39 **Contribution to Fault Detection and Diagnosis: Mixture-based Modelling**
by Tatiana V.Guy, Dani Juričič, Miroslav Kárny and Andrej Rakar, Institute Information Theory and Automation, Academy of Sciences of the Czech Republic/CRCIM
- 40 **Data Mining Applied in Anaerobic Wastewater Treatment**
by Simon Lambert, CCLRC
- 41 **Integrated Multivariate Statistical Process Control and Condition Monitoring for Fed-batch Fermentation Processes**
by Hongwei Zhang, Northeast Wales Institute of Higher Education, UK
- 43 **Saving Energy**
by Claude Lemaréchal, INRIA

R&D AND TECHNOLOGY TRANSFER

- 45 **Solving Complex Problems with Advanced Techniques**
by Enrique Alba, University of Malaga/SpaRCIM
- 47 **Nipper – the Neat Internet Protocol Packet Editor**
Sami Pönkänen and Marko Lyijynen, VTT
- 48 **On Describing B2B Processes with Semantic Web Technologies**
by Santtu Toivonen, Tapio Pitkäranta, VTT and Jung Ung Min, Stanford University, USA
- 50 **Tracking Assets and Vulnerabilities in Corporate Networks**
by Dieter Gantenbein, IBM Zurich Research Lab/SARIT
- 51 **Numerical Simulation through Web Services**
by Stephan Springstubbe, Jürgen Klein and Ottmar Krämer-Fuhrmann, Fraunhofer SCAI
- 53 **Advanced Communication Networks**
by Michel Mandjes, CWI
- 54 **Information Security Research at NTNU**
by Stig F. Mjølsnes and Ingvild Ytrehus, NTNU
- 55 **Supervised Term Weighting for Automated Text Categorization**
by Franca Debole and Fabrizio Sebastiani, ISTI-CNR
- 56 **Remote Monitoring of Health Conditions**
by Dirk Husemann, IBM Zurich Research Lab/SARIT

EVENTS

- 57 **FM 2003 – 12th International Symposium on Formal Methods**
by Stefania Gnesi, ISTI-CNR
- 58 **ERCIM Workshop on Soft Computing 2003**
by Petr Hájek, Institute of Computer Science, Academy of Sciences of the Czech Republic/CRCIM
- 58 **SC 2003 — Industrial Simulation Conference**
by Ana Pajares, Juan C. Guerri and Carlos Palau, Universidad Politécnica de Valencia/SpaRCIM
- 58 **Announcements**

61 IN BRIEF

Cover: Train station in Borlänge, Sweden (photo: Green Cargo). Efficient resource utilisation is a key factor in the railway industry See article by Martin Aronsson on page 27.

Next issue: April 2004
Special theme: Games Technology

In 2004, ERCIM will celebrate its fifteenth anniversary, and it is a time in which to take pride in the ERCIM community's many achievements. ERCIM brings together research institutions from eighteen different countries in Europe, and encompasses more than 12 000 researchers. ERCIM is a de facto premier Network of Excellence in the emerging European Research Area, and a distributed think tank in both information and communications technologies (ICT) and applied mathematics (AM). ERCIM Working Groups are currently active in many significant areas of research within ICT and AM. ERCIM is well positioned to benefit Europe through strategic cooperation with the European Commission, the European Science Foundation, the National Science Foundation in the US, and other international organisations. ERCIM News is an important publication with a wide distribution in both paper and electronic form. ERCIM's high-profile fellowship programme provides young researchers around the world with the opportunity to undertake cutting-edge research at ERCIM member institutions and to become familiar with the diversity of European cultures. Furthermore, the hosting of W3C in Europe is an important activity not only for ERCIM, but also for European industry and the European research community at large. ERCIM is developing quickly, and as it matures, it has the potential to become much more than the sum of its parts. However, ERCIM is currently at a crossroads, and many challenges and opportunities lie ahead. We must meet these challenges and exploit every opportunity in order to enhance Europe's innovation capacity and competitiveness in the global research arena.

ERCIM is currently an organisation of rather loosely coupled members. This coupling is primarily achieved through interactions at Working Group level, joint participation in EU projects and Networks of Excellence, and the dissemination through ERCIM News of R&D results and other developments from member institutions. Nevertheless, cohesion at the overall ERCIM level could be further strengthened. ERCIM must better define its own identity and deliver a clear message to its members, to the broader European research community and to European industry on what may be expected as 'ERCIM added value'. While ERCIM has this issue under constant consideration, it is under some pressure and perhaps feels the need to expand with more members and associate members. However, unless ERCIM first acquires a clearer identity and formulates a strategic plan for growth, innovation and competitiveness, premature and rapid expansion may simply dilute its efforts and make the issue of 'ERCIM added value' more difficult to resolve. In the meantime, ERCIM member institutions ought to be strongly encouraged to promote ERCIM links within the research community of the country they represent.



Stelios Orphanoudakis, Director of the Foundation for Research and Technology - Hellas (FORTH), and President of ERCIM.

Additionally, ERCIM should further develop its own research strategy and act as a catalyst for enhancing the research efforts of its members, while striving to improve and sustain its human resource potential and research infrastructure. ERCIM can become an incubator for the creation and validation of new research ideas in Europe, particularly in interdisciplinary areas where complementarity of skills and ideas is essential. Teams of excellence in selected thematic areas already exist within ERCIM and could be opened to researchers in other academic and research institutions, thus contributing to the creation of long-lasting relationships in European and global research. Given adequate financial support, ERCIM has the potential to play a significant part in a research infrastructure, which is greatly needed if the European Research Area is to become a reality.

This is a difficult period for fundamental research worldwide, and in Europe in particular. We need to ensure that in Europe, with its long-standing tradition in innovative research, fundamental research will remain alive and will be sustained. ERCIM is in a position to play a significant role in this, as it can provide the necessary critical mass and generate much-needed momentum. Fundamental research supports the R&D effort that is needed to sustain the competitiveness of European industry in the long term, and it cannot be allowed to lag behind. ERCIM can serve as an example for cooperation between research institutions and industry in Europe, and in addition should continue to work towards convincing industry, which is also going through a difficult period, that there are benefits in such a long-term cooperation.

ERCIM can refine its strategy for cooperation with the European Commission, the European Science Foundation, the National Science Foundation and other international organisations. In order for such partnerships to be mutually beneficial, it is necessary to define in more detail what ERCIM has to offer, as well as what it can expect. Finally, ERCIM is uniquely positioned to play an important role in technology and research-based education and training, thus contributing to the employability of Europe's workforce in the emerging Information Society. A strategy and a timetable are required here as well, if ERCIM is to take advantage of this opportunity to establish the ERCIM label in the education and training arena at European level.

ERCIM is ready to meet the scientific, organisational, and financial challenges lying ahead and to act as a catalyst for strengthening European excellence in current and emerging research areas.

Stelios Orphanoudakis

Tribute to Gerard van Oortmerssen

On 31 December, 2003, Gerard van Oortmerssen completed his term as President of ERCIM, after five and a half years of distinguished service to the organisation and its member institutions. In May 1998, Gerard van Oortmerssen succeeded Cor Baayen and Dennis Tschritzis to become ERCIM's third president, and was re-elected for a second term in 2001. In November 1999, Gerard hosted and presided over ERCIM's 10th Anniversary celebration in Amsterdam, a memorable event whose success was characterised by a high-level symposium celebrating ERCIM's significant contributions to European research.

During Gerard van Oortmerssen's presidency, ERCIM expanded its membership and geographical coverage to Slovakia, Ireland, Austria, Norway, Luxembourg, Spain and Belgium, with the joining of SRCIM (May 1998), Trinity College (May 2000), AARIT (October 2001), NTNU (January 2002), FNR (July 2002), SpaRCIM (July 2003) and FWO-FNRS (December 2003) respectively. Furthermore, Gerard actively pursued stronger ties with the European Commission, the European Science Foundation and the US National Science Foundation in order to promote ERCIM's role as a distributed think tank in information and communications technology and applied mathematics, as well as a de facto premier Network of Excellence in Europe and throughout the world. In pursuit of the same goal, Gerard adopted a proposal made by the then President of INRIA and Manager of ERCIM, Bernard Larrourou, for ERCIM to become the European host of W3C, which came into effect in January 2003. Within ERCIM, Gerard established the Strategy and Human Resource



Gerard van Oortmerssen was ERCIM's President from May 1998 to December 2003. He has left CWI as managing director as per 1 May 2003. He is now director of TNO Telecom.

Management Taskforces, which have become instrumental in strategic planning within ERCIM and in promoting the sharing of information and experience among ERCIM member institutions on the mobility of researchers and the management of human resources.

Gerard van Oortmerssen leaves ERCIM a stronger organisation and an important player in the emerging European Research Area. Under his leadership, ERCIM has taken important steps towards becoming an open ICT network in Europe and bridging the gap between the research community and industry. This was achieved through its strong relations with startup and other spin-off companies affiliated with its member institutions. Gerard has undoubtedly put his mark on ERCIM, and I am certain that I speak on behalf of the entire ERCIM community when I say he will be missed.

Stelios Orphanoudakis, ERCIM President

Ronald de Wolf — Winner of the 2003 Cor Baayen Award

The 2003 ERCIM Cor Baayen Award was presented to Ronald de Wolf during a ceremony in Luxembourg on 5 November 2003. The 5000 Euro annual award is given to a most promising young researcher having completed a PhD thesis in computer science and applied mathematics, in one of the 'ERCIM countries'.

Ronald de Wolf did his PhD in the area known as quantum information processing and quantum computing. It is remarkable that as an undergraduate he had already written, together with Nienhuys-Cheng, the standard text on a completely different topic called Inductive Logic Programming. This book alone figures as a very impressive PhD thesis. De Wolf received PhD placement offers from MIT and CMU, but decided instead to study at the CWI.

Quantum information processing is a new and important field that combines the physical laws of quantum mechanics with those of computation, resulting in a new model of computation: the quantum computer. The field gained momentum when Peter Shor demonstrated the speed with which such quantum computers could factor large numbers. Most current cryptography, including electronic commerce, is based on the fact that no sufficiently fast factorisation method is known for classical computers. Shor thus

demonstrated that if one could build a quantum computer then most of our cryptographic protocols could be broken.

It is this societal relevance that makes quantum information processing so important. Both North American and European funding agencies have set aside significant amounts of money to invest in this new technology.

Ronald's work received the highest praise in the international community and

during his PhD he received frequent invitations to speak at international meetings and conferences. Among related fields, his PhD thesis ranks among the best worldwide over the last ten years. Ronald's work had a big impact on the field of quantum information processing, and parts of his PhD thesis are now standard inclusions in textbooks on the subject. For example, the impossibility method he developed to show the limitations of quantum computers has become a standard technique and is taught around the world in any class on quantum computing. New protocols and algorithms, such as the quantum fingerprinting technique, have been used by groups such as those at MIT and Berkeley to develop new cryptographic primitives like quantum digital signatures. These advances also attracted significant attention from the press, with articles in *Physics News Update*, *Wired*, *Technology Research News*, *Süddeutsche Zeitung* and others. The



Ronald de Wolf (left) receives the Cor Baayen Award from ERCIM President Gerard van Oortmerssen. On the right is Erna Hennicot-Schoepges, Luxembourg Minister for Culture, Higher Education and Research, who welcomed the participants to the ERCIM meetings in Luxembourg.

work on locally decodable error-correcting codes baffled the best-known researchers in this classical area of computer science. The elegant idea of using quantum mechanics as a way of proving a classical statement is highly original, and was considered a break-

through result in the area of locally decodable codes.

Ronald is one of the very few truly brilliant students that his professor has seen. He is able to combine great mathematical skill with intense creativity in order to solve some of the hardest problems in his field.

The following researchers were nominated for the 2003 Cor Baayen Award: Pekka Abrahamsson, Finland; Giuseppe Amato, Italy; Steffi Beckhaus, Germany; Patricia Bouyer, France; Fredrik Espinoza, Sweden; Andras Gyorgy, Hungary; Keijo Heljanko, Finland; Victor Khomenko, UK; Sebastian Mika, Germany; Lars Rasmusson, Sweden; Zoltan Szabo, Hungary; Gerardo Hermosillo Valadez, France; Ruili Wang, Ireland; Ronald de Wolf, The Netherlands; Albertus P. Zwart, The Netherlands.

More information:

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

Possibilities and Limitations of Quantum Computing

This year's Cor Baayen Award was given to CWI researcher Ronald de Wolf. He received the award for his work on the potential computer of the future: the quantum computer. De Wolf proved several limitations of this revolutionary computation model. He discovered new quantum algorithms that are more efficient than classical algorithms. Furthermore, he successfully applied quantum proof techniques on classical algorithms.

The mantra of quantum computing is the statement 'information is physical'. Information processing and computation are physical processes, and hence subject to the laws of physics. To the best of our knowledge, these laws are quantum mechanical - and quite weird. In contrast, most current computer science is based on models of classical physics. For example, a Turing machine's read-write head is always in a specific position and its memory cells contain definite information (0 or 1). In the 1980s, Richard Feynman and David Deutsch incorporated the laws of quantum mechanics into the model. This gave quite a different picture. In quantum computing the

memory cells can be in a 'superposition' of 0 and 1 at the same time, and different computational paths can 'interfere' with each other destructively or constructively, much like waves.


Roughly speaking, a quantum algorithm works by going through many computational paths simultaneously, in superposition, and using interference to strengthen the paths leading to a correct output. Halfway through the 1990s researchers like Shor and Grover discovered quantum algorithms for factoring integers and searching databases that are more efficient than their classical equivalents. These discoveries made quantum computing one

of the hottest fields around. Whether quantum computers can actually be built is the focus of much research in experimental physics. However, many people feel this effort will bear fruit regardless of the outcome. Even if the construction of a quantum computer is impossible, the research has led to valuable new insights in quantum physics.

The CWI group centred around Paul Vitányi and Harry Buhman was one of the first in Europe to take up quantum computing as a serious research topic. In 1997 de Wolf became a PhD student there and at the University of Amsterdam. Around that time nobody

knew whether quantum computers were superior to their classical equivalents. Together with co-authors, particularly Harry Buhrman, de Wolf proved various strong limitations on quantum computers. For most problems they are not significantly faster than classical computers. These limitations were proved by reducing complexity theoretic questions to algebraic questions about degrees of multivariate polynomials. Sufficient prove of strong -often optimal- lower bounds on the time a quantum algorithm needs to compute a Boolean function, can be given by proving a lower bound on the degree of an n-variate polynomial approximating that function.

Moreover, de Wolf also contributed to the discovery of some quantum algorithms and protocols that outperform their classical counterparts. One example of this is a 'quantum fingerprinting' scheme. It allows two separated parties to compare large chunks of data more efficiently than classical computers. By assigning small quantum states to long classical strings, the amount of data that has to be exchanged for this operation can be exponentially reduced. In the



Quantum computers strengthen the computational paths leading to correct solutions. This can be compared to constructive interference of waves.

future this technique could for example be used to create digital autographs.

After receiving his PhD in September 2001, de Wolf moved to the USA to become a postdoc at UC Berkeley. There he showed, together with Berkeley grad student Jordanis Kerenidis, that quantum computing techniques could be used as a tool to prove new results in classical computer science. Specifically, they analyzed 2-query locally decodable error-correcting codes. This scheme allows each encoded bit to be individu-

ally decoded by looking at only two bits of the codeword, instead of first decoding the entire codeword. Such locally decodable codes have been the focus of much research recently, but the previous best general lower bounds on their length were only polynomial. Using techniques from quantum computing, Kerenidis and de Wolf established for the first time an exponential lower bound. No proof of this result is known that uses just classical proof techniques.

This classical-theorem-with-a-quantum-proof significantly adds to the relevance of quantum computing within computer science in general, since its impact does not depend on the actual building of a quantum computer. As a bonus, the quantum techniques yield quantum protocols for the cryptographic task of 'private information retrieval' that are more efficient than the best known classical protocols. Currently, de Wolf works at CWI as a postdoc. His main research focus is still quantum computing.

Please contact:
 Ronald de Wolf
 Tel: +31205924078
 E-mail: rdewolf@cwi.nl
<http://homepages.cwi.nl/~rdewolf>

Call for Candidates

Cor Baayen Award 2004

The Cor Baayen Award, awarded to a most promising young researcher in computer science and applied mathematics, was created in 1995 to honour the first ERCIM President, is open to any young researcher having completed their PhD thesis in one of the 'ERCIM countries': Austria, Belgium, Czech Republic, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Luxembourg, Norway, Slovakia, Spain, Sweden, Switzerland, The Netherlands and the United Kingdom.

The award consists of a cheque for 5000 Euro together with an award certificate. The selected fellow will be invited to the ERCIM meetings in autumn 2004. A short article on the winner, together with the list of all candidates nominated, will be published in the ERCIM News.

Rules for Nomination

Nominations from each country are made by the corresponding ERCIM Executive Committee member (also referred to as 'national contact'). Those who wish that a particular candidate be nominated should therefore contact the ERCIM Executive Committee member for their country (see www.ercim.org/contacts/execom/). Nominees must have carried out their work in one of the 'ERCIM countries'.

Nominees must have been awarded their PhD (or equivalent) not more than two years prior to the date of nomination. Each ERCIM institute is allowed to nominate up to two candidates from its country. A person can only be nominated once for the Cor Baayen Award. The selection of the Cor Baayen award is the responsibility of the ERCIM Executive Committee.

Submitting a Nomination

To submit a nomination to your national contact, fill out the Cor Baayen Award nomination form (available at the URL given below) and provide a copy of the candidate's PhD thesis as well as copies of the candidate's best papers (max. 5), preferably provided as links to electronic documents.

Deadlines

- Nominations are to be received by the national contacts by 15 April, 2004.
- National contacts are to send their two selected nominations to the coordinator by 30 April, 2004.

Further information can be obtained from your national contact or from the Cor Baayen Award coordinator Lubos Brim, Masaryk University Brno/CRCIM, (lubos.brim@ercim.org).

First ERCIM Workshop on Informatics and Mathematics Applied to Interventional Medicine

by Marc Thiriet

The first workshop of the IM2IM ERCIM Working Group was held from 1-2 December 2003 at CRP Henri Tudor in Luxembourg, as part of the ERCIM biannual meetings hosted by the Luxembourg National Research Fund (FNR).

Participants attended from Canada, Chile, France, Germany, Greece, Italy, Spain, Switzerland and the UK. The set of presentations covered the necessary topics involved in the development of medical simulators.

Medical image processing and virtual reality was discussed by L. Soler, who pointed out the usefulness of navigation tools and medical robotics to our community, a fact which had not hitherto been widely appreciated.

Three-dimensional reconstructions of anatomical organs are a popular research theme (talks from N. Gabrielides, National Technical University of Athens, D. Manoussaki working at FORTH, J. Peiro, Imperial College London and M. Thiriet for the INRIA teams EPIDAURE and REO). Medically oriented flow computations must be performed in image-based discretised domains due to the huge between-subject variability in vessel anatomy and to the patient-dependent shape of the vessel-wall pathologies. Once the three-dimensional reconstruction is made, the mesh must be compatible with the constraints of scientific computing (presentation of L. Baffico, Universidad de Chile and M. Thiriet). Furthermore, mesh adaptation and adaptivity are necessary in the case of unsteady phenomena such as cryotherapy for liver tumours (talk from Youssef Belhamadia, GIREF, Université Laval, Quebec, Canada).

Biofluid flow modelling must also take into account the deformation of the vessel wall during pressure-wave propagation or environment loading, in particular changes in external pressure during respiration, and especially in thoracic vessels. The rheology of the moving fluid and of the bio-conduit wall must be

determined in vivo. At the same time, suitable software must take into account the complete set of rheological properties of the composite bio-tissues, with possible structure changes (presentation of A. Pascau, University of Zaragoza, Spain on Non-Newtonian fluids and M. Thiriet). Due to the geometry of the vessels and the mechanics of the wall and flow genesis, the bioflow models being developed are three-dimensional, generally laminar and quasi-periodic (talk from Giuseppe Pontrelli, CNR, Roma and M. Thiriet). Moreover, although it belongs to a network with the cardiac pump and its input and output impedances, the three-dimensional region of interest is rather limited. Multiphysics multiscale modelling is also being investigated (presentation of M. Fernandez, Ecole Polytechnique Fédérale de Lausanne and J. Peiro). Lumped parameter models may be used for various regions of the vessel network.

New procedures are also being developed however, including the reduced basis element method with fractal and multiphysics homogenisation (talk from Y. Maday, Université Pierre et Marie Curie, Paris). Such complete modelling is associated with a large computational cost. Fast solvers are therefore necessary, using geometric and algebraic multigrid and domain decomposition techniques, and coupling of sequential and parallel codes (presentation of W. Joppich, Fraunhofer-Institute for Algorithms and Scientific Computing (SCAI), Germany).

Mass transport by the fluid conveyed in bio-conduits is investigated in order to understand vessel-wall lesions and to develop new local or bulk therapy strategy. Aerosol transport and deposition in the airways still needs to be accu-

rately modelled in the physical sense (talk from P.-E. Jabin, Ecole Normale Supérieure, Paris). Mini-invasive medical procedures use the body's natural paths, the vessels. Interventional medicine treats vessel-wall damages with medical devices, which are displaced through the vascular bed and deployed within the lesion using catheter-based techniques. The design and shape optimisation of such devices, as well as ventricle-assist pumps is currently of major interest (presentation of M.-I. Farinas, Ecole Polytechnique de Montréal, Canada).

During the final session it was decided that the next IM2IM workshop would be held in Marseille in August 2004, in order to coincide with a six-week summer school in biomathematics. It is planned to hold the following workshop at IRCAD in Strasbourg in December 2005. The Working Group intends to submit a proposal under the European Commission's Future and Emerging Technologies scheme to drive the Working Group in a precise direction.

Link:

<http://www-rocq1.inria.fr/Marc.Thiriet/Im2im/>

Please contact:

Marc Thiriet, INRIA

E-mail: marc.thiriet@ercim.org

XForms: the Next Generation of Web Forms

by Steven Pemberton, CWI and W3C

HTML Forms formed the basis of the e-commerce revolution. Now after 10 years experience with them, W3C has released XForms, the next generation of Web Forms. In contrast to HTML forms, in which functional and presentation markup are intertwined, XForms lets forms authors distinguish the descriptions of the purpose of the form, the presentation of the form, and how the results (the instance data) are written in XML.

The design of XForms is based on an analysis of how HTML Forms have been used in the last ten years, and what could be changed to make them even more useful. In brief, XForms lets you do everything that HTML Forms can do, and more, such as:

- check data values while the user is typing them in
- indicate that certain fields are required, and that the form cannot be submitted without them
- constrain values in certain ways, such as requiring them to be in a certain range
- submit forms data as XML
- integrate with Web services, for instance by using SOAP and XML RPC
- submit the same form to different servers (for instance a search string to different search engines)
- save and restore a form to and from a file
- use the result of a submit as input to a further form
- get the initial data for a form from an external document
- calculate submitted values from other values
- build 'shopping basket' and 'wizard' style forms without needing to resort to scripting.

Properties of XForms

In detail, XForms has several important properties:

- *XForms improves the user experience:* XForms has been designed to allow much to be checked by the browser, such as types of fields being filled in, that a particular field is required, or that one date is later than another. This reduces the need for round trips to the server or for extensive script-based solutions, and improves the user experience by giving immediate feedback on what is being filled in.
- *It makes it easier to author complicated forms:* Because XForms uses declarative markup to declare properties of values, and to build relationships between values, it is much easier for the author to create complicated, adaptive forms, without having to resort to scripting. An HTML Form converted to XForms looks pretty much the same, but when you start to build forms that HTML wasn't designed for, XForms becomes much simpler. See XForms for HTML Authors for some examples.
- *It is XML, and it can submit XML:* XForms is properly integrated into XML: it is in XML, the data it collects in the form is XML, it can load external XML documents as initial data, and can submit the results as XML. By including the user in the XML pipeline, it at last means you can have end-to-end XML, right up to the user's desktop.

- *It combines existing XML technologies:* Rather than reinventing the wheel, XForms uses a number of existing XML technologies, such as XPath for addressing and calculating values, and XML Schema for defining data types. This has a dual benefit: ease of learning for people who already know these technologies, and the ability for implementors to use off-the-shelf components to build their systems.
- *It is device independent:* The same form can be delivered *without* change to a traditional browser, a PDA, a mobile phone, a voice browser, and even some more exotic emerging clients such as an Instant Messenger. This greatly eases providing forms to a wide audience, since forms only need to be authored once.
- *It is internationalized:* Because the data submitted is XML, it is properly internationalized.
- *It is accessible:* XForms has been designed so that it will work as well with accessible technologies (for instance for blind users) as with traditional visual browsers.

Implementations

The XForms Working Group that created XForms contains representatives from many major computing companies. On top of this there are a large number of implementations emerging. In fact XForms is the most-implemented W3C specification ever at this stage in its life-cycle.

There are several implementations that let you use XForms with existing browsers, either with plugins or by transforming the forms at the server, as well as a number of browsers that implement XForms natively.

At the time of writing, XForms has been demonstrated on PDAs and mobile phones using proxies, using the same techniques that are often used to present web content on smaller devices. We anticipate that native XForms clients will appear on mobile devices soon.

There is a list of XForms Implementations that contains details of many of these implementations; an editorial review including screen shots of several of them can be found at www.xml.com/pub/a/2003/09/10/xforms.html. XForms does not need any extra support from servers: it works with all current servers. Although XForms has several new submission methods compared with HTML Forms, these are all standard HTTP methods for sending data over the network. XForms can use 'PUT' to put data to a server, can talk to a SOAP or XML RPC servers, but also supports legacy formats allowing forms to work with all existing forms servers.

Finally, some large user populations are beginning to emerge, including some government e-government initiatives, a country's tax service, and a country's insurance industry.

Steven Pemberton is a senior researcher at the CWI in Amsterdam. He is chair of the W3C XForms Working Group.

Links:

XForms home page: <http://www.w3.org/Markup/Forms>

Introduction to XForms for HTML Authors:

<http://www.w3.org/Markup/Forms/2003/xforms-for-html-authors.htm>

XForms implementations:

<http://www.w3.org/Markup/Forms/#implementations>

W3C Launched Spanish Office

W3C is pleased to announce the opening the W3C Spanish Office, based at the Fundación para el Fomento de la Investigación Científica y la Tecnología (FICYT) in Oviedo. Although the primary goal of this Office is to be dedicated to outreach in Spain, this is also the first W3C Office with active contacts with Spanish-speakers worldwide. The new W3C Spanish Office will be of help in coordinating with Latin American communities by also disseminating information in Latin America, encouraging and co-ordinating new translations to Spanish, etc.

The W3C Spanish Office joins the complement of W3C European Offices in The Benelux Countries, Finland, Germany & Austria; Greece; Hungary, Italy; Sweden; and the United Kingdom and Ireland. Francisco Sanchez is Office Manager. The opening ceremonies were held on 20 October in Oviedo.

Links:

W3C Spanish Office Web site: <http://www.w3c-es.org/>
 Press release: <http://www.w3.org/2003/10/spoffice-pressrelease>

W3C Requests Eolas Patent Review

On 29 October the W3C has presented the United States Patent and Trademark Office with prior art establishing that US Patent No. 5,838,906 (the '906 patent) is invalid and should therefore be re-examined in order to eliminate this unjustified impediment to the operation of the Web, because the critical prior art was neither considered at the time the patent was initially examined and granted, nor during recent patent infringement litigation.

In an unprecedented step, Tim Berners-Lee, W3C Director and inventor of the Web, sent a letter to Under Secretary Rogan requesting that his office reinvestigate the matter. "W3C urges the USPTO to initiate a reexamination of the '906 patent in order to prevent substantial economic and technical damage to the operation of World Wide Web," stated Berners-Lee. "The impact of this patent will be felt not only by those who are alleged to directly infringe, but all whose web pages and application rely on the stable, standards-based operation of browsers threatened by this patent. In many cases, those who will be forced to incur the cost of modifying Web pages or software applications do not even themselves infringe the patent - assuming it is even valid."

US Patent 5,838,906 'Distributed hypermedia method for automatically invoking external application providing interaction and display of embedded objects within a hypermedia document' was filed in October 1994 and granted in November 1998. The patent, owned by the University of California with licensing handled through Eolas Technologies, covers mechanisms for embedding objects within distributed hypermedia documents, where at least some of the object's data is located external to the document, and there is a control path to the object's implementation to

support user interaction with the object. The implementation can be local or distributed across a network, and is automatically invoked based upon type information in the document or associated with the object's data.

The practical impact of withholding unrestricted access to the patented technology from use by the Web community will be to substantially impair the usability of the Web for hundreds of millions of individuals in the United States and around the world. Nearly every Web user relies on plug-in applications that add services such as streaming audio and video, advanced graphics and a variety of special purpose tools. On 30 October, the US PTO have announced that they will be conducting a re-examination of all claims relating to US Patent 5,838,906 citing "a substantial outcry from a widespread segment of the affected industry."

Links:

20 Oct. 03 press release: <http://www.w3.org/2003/10/28-906-briefing>
 TimBerners-Lee's letter: <http://www.w3.org/2003/10/27-rogan.html>
 Prior art filing: <http://www.w3.org/2003/10/301-filing.html>
 FAQ: <http://www.w3.org/2003/09/public-faq.html>
 W3C Holds Ad Hoc Meeting on Recent Court Decision, Launches Public Discussion List: <http://www.w3.org/2003/08/patent>

W3C Held its First Outreach Event in Mainland China

The China International Forum on WWW's Development 2003 was held in Beijing on 12-13 November. Attendees discussed the future of the Web, accessibility, SVG, the mobile Web, the Semantic Web and internationalization. The event was co-organized by the China Computer Federation and the W3C Office in Hong Kong.

Links:

China International Forum on WWW's Development 2003 Web site: <http://www.w3c.org.hk/CIFWeb03/index.html>
 Press Release: <http://www.w3.org/2003/10/beijingevent-pressrelease>

Last W3C Recommendations

- Portable Network Graphics (PNG) Specification (Second Edition): First published 1 October 1996, revised 10 November 2003, David Duce
- XForms 1.0: 14 October 2003, Micah Dubinko, Leigh L. Klotz, Roland Merrick, T. V. Raman
- XML Events: 14 October 2003, Shane McCarron, Steven Pemberton, T. V. Raman

An exhaustive list of all W3C Technical Reports:

<http://www.w3.org/TR/>

ERCIM is the European host of W3C.

Industrial Diagnosis, Planning and Simulation — Introduction

by Per Kreuger

The special theme of industrial diagnosis, planning and simulation is structured around a selection of industrial problems rather than a particular type of technology or methodology. We have chosen contributions describing a variety of techniques and applications that are of relevance to these problems, with the intent of highlighting state-of-the-art applied research within the ERCIM institutes and their associates.

Recently at SICS, the Industrial Applications (IAM) Laboratory was formed around a novel topic in applied research. We brought together a number of researchers with expertise in mathematics and computer science, who share the conviction that the study of the applicability of techniques and methods developed within the scientific community should itself be the subject of research. The activities of such a group will, of course, always be influenced by the backgrounds of the people participating in it, not only with regard to their specific scientific expertise but also, and in this case significantly, by their familiarity with practical problems outside the scope of their scientific schooling. The group formed at SICS brings together researchers from the fields of combinatorial optimisation, scheduling, flow theory, constraint programming, statistical methods for data analysis and learning systems. They have domain expertise in process and manufacturing industry and resource management, monitoring and optimisation in transportation and logistics.

The initiative behind the formation of this group came from the observation that methods and techniques from these scientific disciplines often have the potential to improve both the efficiency of resource usage and the quality of products and services, as well as reducing the costs associated with production. On the other hand, not all methods proposed by researchers scale well to the size and diversity of industrial problems.

The fact that techniques and methods developed in science are used to solve practical problems is, of course, nothing new. However, while this is the basis of all technology, the quality of the interplay between science and technology within a society determines to a large extent its potential for growth and development. In the case of computing science, which is the study of algorithms, their properties and use, one must always be careful to verify results against practical problems. This is a time-consuming and, at times, frustrating experience for many researchers, since their experience lies in a different field to that in which the problem occurs. Even so, this process is essential

for the growth of technology and provides valuable feedback to the research discipline.

For example, within the field of optimisation and combinatorial reasoning there exist many elegant and highly applicable results for a wide selection of specific problems. In a few cases, an industrial problem will exactly match a certain problem idealisation and the mathematical model used to study the properties of a particular class of algorithms. This is, however, only rarely the case. Real practical problems in, for example, industrial production, are invariably complex combinations of several smaller problems and may have numerous related idiosyncratic conditions. The ability to understand and describe a practical problem in terms of the types of models and methods used in this field is usually referred to as ‘modelling expertise’. Interestingly, this expertise is rarely described in the scientific literature. The ability to use outside a scientific field the algorithmic methods developed within that field exists in parallel with the science, and is indeed rather more of an artform.

A systematic method for solving a class of practical problems is called a methodology. However, most methodologies used in computing science are based around a fairly limited set of algorithmic methods and quickly become useless whenever the problem changes too much or is combined with other problems into a more complex situation. This raises issues of method generality and scalability, which should be of prime concern to the scientific field developing the method as well as of great practical importance to the society in which the scientist works. A systematic study of these issues will lead to two types of results:

- better understanding of typical models for a selection of important real-life industrial problems
- better understanding of the properties of a selection of practically useful algorithmic methods.

Industrial applications of state-of-the-art techniques and methods in computing science allow researchers to test, in practice, the flexibility, scalability and utility of their techniques and methods. In addition, it is an opportunity to push the mature parts of the scientific field out into practical use for the benefit of the industry and society at large, be it in manufacturing, transportation, processing, telecommunications, biotechnology or the service industry. This is true for many fields of computer science, but particularly for those directly involved in the modelling and solving of typical problems occurring in industry, eg process simulation, monitoring and prediction, capacity analysis and allocation, fault detection and

diagnosis, production planning, flow optimisation, resource scheduling and allocation, structure detection and matching. We believe there is a strong and fundamental need to systematically study the practical utility of such methods by applying them to a large number of real cases.

For the special theme we have aimed to highlight successful or promising applications of advanced techniques from these fields, including some in full-scale industrial settings, which we hope will contribute to our understanding of the strengths and weaknesses of the various techniques and methods. The articles have been grouped into four sections:

1. Electronics and Networks
2. Process Monitoring and Optimisation
3. Transportation and Logistics
4. Process Design and Management.

The first section contains two articles on the use of magnetic field models in the electronics industry, one on the management and optimisation of communications network resources, and a fourth on intrusion prevention.

The second – and largest – contains articles on monitoring, diagnosing and optimising production processes. Of these, several are concerned with the detection of deviations from normal process parameters, while others address issues in connection with identifying the cause of, and correcting, faulty behaviour. In a few cases, optimisation of process parameters and redundancy of monitoring systems are also relevant.

The third section contains articles on problems in transportation and logistics. This type of issue often contains instances of specific problems that are comparatively well understood, but in practice are frequently made very difficult to solve by combinations of sub-problems such as resource allocation, routing and scheduling. One of the articles describes an analytical method for capacity analysis and assessment using techniques from discrete event systems. Three are concerned with management and routing of vehicles, which is one of the most cost-intensive operations in transportation. Others describe approaches to solving packing, placement and storage problems.

The articles in the final section describe methods for process design and project management — enjoy!

Link:

<http://www.sics.se/iam>

Please contact:

Per Kreuger, SICS, E-mail: piak@sics.se

ARTICLES IN THIS SECTION

- 10 Industrial Diagnosis, Planning and Simulation - Introduction**
by Per Kreuger, SICS
- 12 Efficient Algorithms for Inverse Magnetic Field Problems**
by Domenico Lahaye, CWI
- 13 Deviation Detection of Industrial Processes**
by Anders Holst, Jan Ekman and Daniel Gillblad, SICS
- 14 Bidding Agents Acquire Cargo in Online Auctions**
by Pieter Jan 't Hoen, CWI
- 15 Decision Support System for Process Analysis and Supervision**
by Galia Weidl, University of Stuttgart, Germany
- 17 Mathematical Challenges in the Electronics Industry**
by Wil Schilders, Eindhoven University of Technology, The Netherlands
- 18 Contpack - Maximising the Volume Utilisation of Containers**
by Onno Garms, Ralf Heckmann and Stefan Rank, Fraunhofer SCAI
- 20 First Commercial Bayesian Software for Intelligent Troubleshooting and Diagnostics**
by Claus Skaanning, Dezide, Denmark
- 22 Innovating Shoe Manufacturing using Advanced Simulation Techniques**
by Emanuele Carpanzano and Andrea Cataldo, ITIA-CNR
- 23 Billy Goat Detects Worms and Viruses**
by James Riordan and Diego Zamboni, IBM Zurich Research Lab/SARIT
- 24 Decision Support for Personnel and Resources Management in Multi-Project Environments**
by Antonia Bertolino, Eda Marchetti and Raffaella Mirandola, ISTI-CNR
- 26 Towards Useful Tools for Estimating Railway Infrastructure Capacity**
by Jan Ekman, SICS
- 27 Efficient Utilisation of Resources in the Railway Industry by Levelling Departures and Arrivals of Trains**
by Martin Aronsson, SICS
- 28 Real-Time Optimisation of Empty Rail Freight Car Distribution**
by Martin Joborn, Carmen Consulting, Sweden
- 29 KUORMA: A Collection of APS-Algorithms for Forest Industry Wood Transport**
by Juha Savola, VTT
- 31 Parc Bandwidth on Demand**
by Helmut Simonis, Parc Technologies Ltd, UK
- 32 DesParO - A Design Parameter Optimisation Toolbox using an Iterative Kriging Algorithm**
by Christof Bäuerle, Clemens-August Thole and Ulrich Trottenberg, Fraunhofer SCAI
- 33 On-line Testing of the Reactor Protection System in the Paks Nuclear Power Plant**
by István Varga, Tamás Bartha and Alexandros Soumelidis, SZTAKI
- 35 Process Systems: Theory and Applications from Different Aspects**
by Katalin Hangos and Gábor Szederkényi, SZTAKI
- 37 Multi-Agent System Technology in a Port Container Terminal Automation**
by Vicent J. Botti, Universidad Politécnic de Valencia/SpaRCIM
- 39 Contribution to Fault Detection and Diagnosis: Mixture-based Modelling**
by Tatiana V.Guy, Dani Juričić, Miroslav Kárny and Andrej Rakar, Institute Information Theory and Automation, Academy of Sciences of the Czech Republic/CRCIM
- 40 Data Mining Applied in Anaerobic Wastewater Treatment**
by Simon Lambert, CCLRC
- 41 Integrated Multivariate Statistical Process Control and Condition Monitoring for Fed-batch Fermentation Processes**
by Hongwei Zhang, Northeast Wales Institute of Higher Education, UK
- 43 Saving Energy**
by Claude Lemaréchal, INRIA

Efficient Algorithms for Inverse Magnetic Field Problems

by Domenico Lahaye

The design of electromagnetic devices can be stated in mathematical terms as inverse magnetic field problems. Solving these hard optimisation problems efficiently is of prime industrial relevance. At CWI, research has started on a new class of algorithms for inverse problems based on the space-mapping algorithm. This algorithm exploits a multiscale representation of the space of design variables, in order to speed up the convergence of the optimisation process. Applications include the design of low-frequency electromagnetic energy transducers.

In recent decades, the use of mathematical modelling and numerical techniques in the industrial design of electrical machines and transformers has gained popularity. Mathematical models for electromagnetic energy transducers are built starting from the Maxwell equations - a set of coupled partial

differential equations describing electric and magnetic fields. These equations are treated numerically by discretising them, for example, by the finite element method.



In industrial design practice, so-called inverse magnetic field problems arise frequently. These are optimisation problems in which one seeks a source distribution that produces an a priori specified magnetic field. These hard, nonlinear optimisation problems are mathematically intriguing due to their ill-posedness and non-unique solutions. Furthermore, finding solutions efficiently is critical, and is a key ingredient in the design of high-gain electromagnetic energy transducers that are less susceptible to undesirable parasitic effects such as heat, noise and mechanical vibrations.

These inverse magnetic problems include design variables such as descriptions of the geometry of a device (eg the width of the air gap in an electrical machine), and external loading conditions in terms of applied currents and voltages (eg the current through the coils in an MRI scanning instrument). The evaluation of the function to be minimised (the cost function) for given values of the design variables requires a finite element simulation. These simulations form a computational bottleneck,

especially in performing full three-dimensional or transient modelling. Recently, the MAS2 research group at CWI started a research project with the aim of improving existing solution techniques for inverse magnetic field problems. The goal of this research project is to construct and apply new optimisation algorithms that reduce the required number of cost function evaluations. These algorithms are based on a recently introduced idea known as 'space mapping'.

Space-mapping algorithms exploit a multiscale representation of the search

space in order to accelerate the convergence of the iterative optimisation process. While fine-scale models correspond to the previously mentioned finite element models, coarse-scale models are simple analytic models for the device under consideration. Examples include magnetic equivalent circuits (equivalent to electrical circuits) or analytical expressions for the electromagnetic force and torque (such as the Biot-Savart law). The coarse-scale representation allows well-established rules of thumb to be incorporated in the optimisation process. The mapping from fine to coarse space is performed through a



For current-carrying coil and a permanent magnet, the first figure (left) shows the equipotential distribution computed by the finite element method.

The second figure shows the triangular mesh used for the computation.

parameter extraction process. The space-mapping technique was introduced in the context of microwave engineering.

In our project, we must solve quasi-stationary inverse magnetic field problems. The aims of the project are twofold. Firstly, we want to provide a theoretical framework for the space-mapping technique in order to gain insight into its convergence properties. We will do so by establishing links between the space-mapping technique and seemingly closely related defect correction techniques. For the latter, a sound mathematical theory exists. Secondly, we want to apply the space-mapping technique for the solution of quasi-stationary inverse magnetic field problems. A typical application that we have in mind is the design of a linear

actuator to be used in a high-accuracy positioning system. This research effort is a collaboration between the MAS2 research group and the Electromechanics and Power Electronics (EPE) research group at the Eindhoven University of Technology, who will provide their expertise in modelling and design of electromagnetic devices and in the validation of design through experiments. The Dutch Ministry of Economic Affairs is funding the project through an intensive research program in electromagnetic power conversion technology (IOP-EMVT).

We envisage two possible extensions of the current project. In the first of these, we plan to apply the space technique in the design of contactless energy transfer systems. In these systems, electrical

energy is transmitted by magnetic induction, and electrical connections and physical contact are avoided. Recent applications of CETS have emerged, for example, in the transfer of energy to mobile platforms, in the intervention-free charging of household appliances and in medical scanning systems. In the second project extension, we plan to solve inverse problems in which magnetic and thermal aspects are coupled. Solving these problems is an important step in the design of electromagnetic energy transformers that minimise heat loss.

Link:

<http://www.cwi.nl/mas2>

Please contact:

Domenico Lahaye, CWI

Tel: +31 20 592 4122

E-mail: Domenico.Lahaye@cwi.nl

Deviation Detection of Industrial Processes

by Anders Holst, Jan Ekman and Daniel Gillblad

One of the major problems in complex processes is that it is impossible to foresee every possible type of fault or abnormal situation. This makes it hard to construct automatic fault detection systems, since it is impossible to know what to look for. An interesting approach to this problem is deviation detection, which turns the problem around by using learning systems to model 'normal' behaviour and then measuring the abnormality of a new situation against this.


The general approach for deviation detection is to construct a statistical model of the relevant features of the domain being monitored. The parameters of the statistical model are estimated using data consisting of (only or predominantly) normal cases. When used for detecting deviations, the model calculates the likelihood of a new sample case, ie, the probability that the model would generate that sample. A low likelihood signals that the sample is not likely to be drawn from the model of the normal data, and thus may be abnormal.

The area where this approach has received the most attention is in intrusion detection systems (IDS). Here, statistical methods are used to characterise traffic or user behaviour in various ways, based on variables such as traffic volume, distribution over time, types of accessed

ERCIM News No. 56, January 2004

documents, multiple accesses to different servers etc. At SICS we have used a similar approach to create a system to detect fraud in ADSL-based pay-per-view systems in a project with TeliaSonera. A number of features from the domain, such as the number of ordered and billed films and the amount

of outgoing and incoming network traffic, was used to train a mixture model to recognise normal (non-fraudulent) users. When tested on another set of normal users, plus a number of fraudulent cases of three different types (billing fraud, subscription fraud, and illegal redistribution fraud), the model success-



A projection of the fraud detection data. The redness of the points indicates how unlikely they are (the negative log likelihood).

fully picked almost all of the billing and redistribution frauds, and the majority of the subscription frauds.

Another application in which SICS is applying the same techniques is the detection of abnormal behaviour in ship movements based on radar images. The goal is to be able to generate a warning if a ship is moving in a strange, dangerous, or suspicious way. The relevant features that are modelled in this case include local properties, such as current position and speed, as well as more qualitative features of the route travelled so far, such as the number of turns, stops, possible encounters with other boats, and so on.

It is also becoming possible to embed such normality modelling into industrial systems such as network routers, industrial robots, the components of paper mills etc. This will make it possible for the components themselves to generate a

warning when something is beginning to get out of order; for instance, in the context of condition-based maintenance, the components could indicate when they are close to wearing out and need replacement or service. Making components self-monitoring in this way has several advantages in terms of reliability, uptime, reduced waste of materials, and lower maintenance and production costs.

We are taking the first steps towards this in a project funded by SSF (Swedish Foundation for Strategic Research) and entitled 'An 'English butler' for the process industry'. One of the test cases is a Swedish hot steel rolling mill. There are many factors affecting the final quality of the steel, such as the form of the heating curve. By monitoring these factors, it is possible to detect when the process moves outside the normal region in state space, and hence when a risk of damage to the steel exists.

Critical for the success of deviation detection is the use of sufficiently powerful statistical learning models to catch relevant behaviour in often highly complex domains. At SICS we have developed a general statistical modelling tool, 'Hierarchical Graph Mixtures', based on a combination of mixture models, probabilistic graph models, and Bayesian statistics. This has given us a very powerful and flexible framework for modelling, which is applicable to a wide range of real-world tasks.

Links:

- <http://www.sics.se/iam/projects/butler.html>
- <http://www.sics.se/iam/projects/adaf.html>
- <http://www.sics.se/iam/projects/dallas.html>

Please contact:

Anders Holst, Jan Ekman, Daniel Gillblad, SICS
 E-mail: aho@sics.se, jan@sics.se, dgi@sics.se

Bidding Agents Acquire Cargo in Online Auctions

by Pieter Jan 't Hoen

A growing field of applications for Multi-Agent Systems is the world of logistics and planning. CWI's research in the DEAL project focuses on optimising the use of road transport capacity through acquisition of cargo by agents bidding in online auctions. This promises a substantial gain in profit.

CWI's Evolutionary Systems and Applied Algorithms group headed by Prof. J.A. La Poutré is a participant in the Distributed Engine for Advanced Logistics (DEAL) project. The goal of this project is to optimise the use of road transport capacity in a logistics setting, using agent-based technology. Among other things, this involves providing an open market where business participants can conduct online electronic acquisitions and place orders.

As real-world scenarios in logistics and planning grow in size and complexity, classical centralised approaches become unwieldy and brittle. In a logistics setting, a centralised planner can be overwhelmed by the stochastic nature of real-world events. For example, new business opportunities arise that cannot be effi-

ciently exploited. Furthermore, schedules can go awry and their repair can require extensive rescheduling. A Multi-Agent System (MAS) offers a robust alternative to a centralised planning approach. An individual agent within the system is an autonomous and adaptive piece of software. It is therefore able to learn to respond effectively to changing scenarios, and can efficiently solve tasks requiring multiple agents by communicating with and coordinating other agents. For example, MASs have been successfully applied to difficult supply-chain management problems where centralised control has failed. In the transportation setting, agents, which act for individual trucks or coalitions of transporters, are able to continuously optimise and update planning as incidents occur and new opportunities for

profit arise. An agent can quickly assess the value of new cargo or inform relevant agents if a planned delivery is delayed because of traffic or an unexpected overhead in the unloading of a truck.

For application to transportation settings, research at CWI focuses on possible market mechanisms and bidding and negotiation strategies of the agents. CWI is developing robust, distributed market mechanisms as part of MASs for usage in (simulations of) the transportation sector. Online decentralised auctions serve as the model, where agents - representing trucks - bid for cargo. Current centralised planning is only able to utilise 40 to 60% of the total truck haulage capacity. Preliminary experiments show that continuous acquisition of new orders by agents bidding in

online auctions (spot markets) can raise this to 80%.

A bidding strategy under further development at CWI is the use of decommitment as a tool for handling unknown future orders available for auction. Using a decommitment strategy, an agent is able to postpone the transportation of an acquired load to a more opportune time. This allows the agent/truck to bid in auctions it would not otherwise consider due to prior commitments. Experiments show that substantial increases in performance through the use of a decommitment strategy can be expected due to additional flexibility in the planning capabilities of the MAS. This increase in performance for the (abstract) model of the transportation sector can be seen as a lower bound for expected additional performance in practice. This claim is substantiated through experiments (see the links below for publications), which have shown that the relative impact of a decommitment strategy increases with the complexity of the world.

A full-scale MAS also facilitates communication and therefore advanced negotiation between agents. Logistics planning is faced with complex constraints and the



Agents increase profits by continuously acquiring new cargo for transportation in online auctions.

varying goals and capabilities of the transporters. CWI has demonstrated through evolutionary simulation that agents can learn to negotiate successfully over complex multi-issue goals (see the links below for publications). Agents can also learn to exchange tasks/orders so as to improve performance for both parties. Negotiations between agents allow for continuous optimisation, while adaptivity promotes efficient negotiation strategies with specific opponents. Furthermore, agents can negotiate to form coalitions for the execution of complex tasks. For example, an order too large to be transported by one truck can be split up over a

temporarily cooperating group of agents.

The DEAL project intends to deliver a working prototype of a platform supporting online acquisition of cargo using MASs in three years' time. Research will continue at CWI into the precise conditions and parameters required in the underlying model of decentralised auctions for successful application of an MAS in such a logistics setting. The research focuses on the market mechanisms required, ie

which rules guide the bidding behaviour of the agents and how can agents best negotiate over a transport? This is comparable to the study of rules as found in a market economy. The application of the right type of intelligent software will play a major role.

Links:

<http://www.almende.com/deal>
<http://www.cwi.nl/sen4>
<http://www.cwi.nl/~hoen>

Please contact:

Pieter Jan 't Hoen, CWI
 Tel: +31 20 592 4276
 E-mail: hoen@cwi.nl

Decision Support System for Process Analysis and Supervision

by Galia Weidl

The increasing complexity of industrial processes and the struggle for cost reduction and higher profitability means automated systems for processes diagnosis in plant operation and maintenance are required. The system proposed here is designed to address this issue.

As system complexity increases, condition monitoring and fault diagnosis become demanding tasks for process operators, who face an overflow of data and may have to undertake extensive searches for unexpected faults. This can lead to work overload and high stress levels. Simultaneously, technical, economical and time constraints on production efficiency and quality require that condition monitoring and adaptation of production and service processes be undertaken on a continuous basis.

In this context the ABB group, along with its customers from process industries and the Danish company Hugin Expert, have collaborated on a project (which commenced in 2000) targeting the development of a methodology for root cause analysis (RCA) and decision support on process operation. The author, who was involved in technology evaluation and methodology development while associated with ABB, is now associated with IFF – Stuttgart University and is working on an exten-

sion of the methodology to automotive and manufacturing industries and services. It is expected that this extension will be implemented within a European project known as 'The Butler Concept' (see page 16 in this issue).

Motivation and the Techniques Employed

In finding the origin of a process disturbance or fault, there exists a need for a quick and flexible guidance tool for decision support at higher automation levels,

including analysis of process conditions and advice on cost-efficient actions.

The technology of probabilistic graphical models such as Bayesian networks and influence diagrams have turned out to be the best of a number of alternatives, in the case where high diagnostics capabilities, explanation of conclusions for transparency in reasoning, and trustworthy decision support are expected by the users (process engineers, operators and maintenance crew).

Due to the existence of a number of first-level diagnostic tools, the aim has been to provide decision support in process operation. The framework of Bayesian networks has been found to be an efficient and flexible tool in overall-level process operation analysis, since not all conditions are measurable or computable in real time, and the combinatorial reasoning procedure is subject to uncertainties.

The development of the methodology incorporated the following system requirements and modelling issues:

- root cause analysis of industrial processes with adaptation to process operation/grade changes, aging and wear
- reusable system design for various process applications
- reusable modelling of repetitive structures (eg sensors, control loops) and assets (pumps, valves)
- risk assessment of disturbances by analysis of signals' level-trend, while

The Butler Concept

This article by Galia Weidl as well as the articles 'Deviation Detection of Industrial Processes' on page 14 and 'First Commercial Bayesian Software for Intelligent Troubleshooting and Diagnostics' on page 20 are related to the 'Embedded Butler Concept'. The idea is to create an embedded and self-learning system for detecting, diagnosing, and correcting faults and deviations in complex systems.

A consortium containing among others SICS (coordinator), ABB, Aalborg University, University of Stuttgart, Tampere University of Technology, Mälardalen University, Warsaw University of Technology, Hugin Expert A/S, and Dezide ApS, has been formed and an application for a STREP project has been submitted under the Embedded Systems call in the 6th Framework Programme of the European Commission. The goal of the project is to develop advanced functionality within each of the sub-components of fault detection, fault diagnosis, and corrective actions, and to combine that functionality into a complete system.

<http://www.sics.se/iam/>

Please contact: Björn Levin, SICS, blevin@sics.se

- adapting to changes in process operation mode
- ease of communication and explanations of conclusions at different process levels, eg process overview, (sub)sections, units and instrumentation.

These modelling and system requirements have been met in the methodology. This is supported by the integration of the methodology and the Hugin-tool into the ABB Industrial IT platform. This RCA-integration allows efficient data exchange with all available IT-applications, eg distributed control systems, diagnostics of sensors and control loops, and physical-model computations. The infrastructure for applying this methodology in different domains is therefore ready for immediate use. The method and the analyser for producing information have been the

subject of five separate patent applications.

The first prototype of the root cause analysis system was tested by ABB with real process data during 2002. The monitoring and root cause analysis of the digester operating conditions in a pulp plant have been chosen as a real-world application (see Figure 1). The structure of one of the developed Bayesian Networks is shown in Figure 2.

The real application development has been closely related with its integration into the Industrial IT platform. The extension includes the use of object-oriented Bayesian networks (OOBN). OOBN facilitate the modelling of large and complex domains and allow reusability.

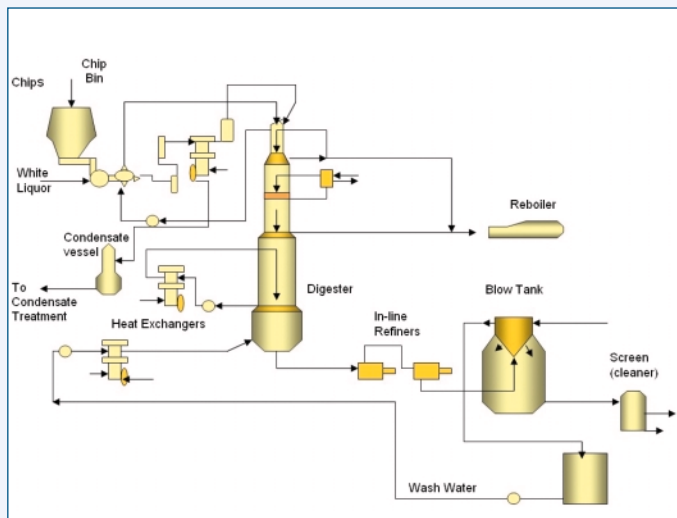


Figure 1: Digester Fiber-line. Case-study: Monitoring of the digester operating conditions.

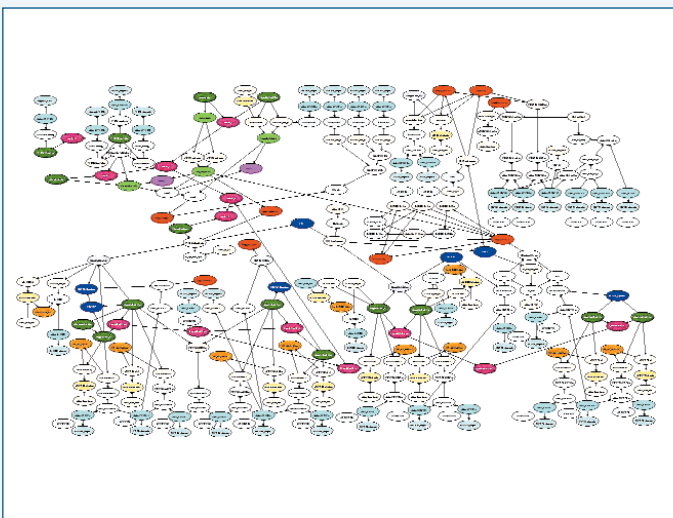


Figure 2: An example of a Bayesian Network for root cause analysis of process operation.

Benefits, Results and Consequences

The resulting RCA-system provides process operators with information on condition overview and advice on the most efficient sequence of corrective actions. In addition, the advantages for process industries include the seamless integration of all relevant information sources on the Industrial IT platform, resulting in automated and early assessment of abnormal conditions with flexible diagnosis and advice. Future enhancements will involve integration of the functionalities providing advice on a suitable time for maintenance activities under technical constraints and order deadlines in process scheduling, as well as simulation of the impact of intended

corrective actions on process efficiency. This allows pro-active (instead of reactive) troubleshooting to be undertaken, which increases the process performance, availability and output, avoids potential process breakdowns and cuts both operation downtime and maintenance cost.

**Expected Function –
RCA as a Powerful Complement
to System Control**

We believe that this methodology will provide a powerful complement to system control in process industries. Our expectations and the proof of the system's potential are based on successful tests at ABB with real plant

data. Future challenges are primarily related to the application of hybrid learning systems for abnormality detection, diagnosis and advice in large-scale industrial settings.

Links:

<http://www.abb.com>
<http://www.hugin.com>
<http://www.iff.uni-stuttgart.de>

Press-release Web Site:

<http://www.hugin.com/cases/Industry/ABB/ABB18122003.pdf>

Please contact:

Galia Weidl, University Stuttgart
 E-mail: giw@iff.uni-stuttgart.de,
galia@mathematik.uni-stuttgart.de

Mathematical Challenges in the Electronics Industry

by Wil Schilders

New generations of electronic devices and integrated circuits are characterised by shrinking dimensions and increasing frequencies. Electromagnetic effects that were previously neglected must now be taken into account in order to reliably predict behaviour. The Scientific Computing Group of the TU Eindhoven is developing new numerical techniques for coping with this challenging problem.

The electronics industry has always been a very fertile working environment for mathematicians. In the 80s, the simulation of semiconductor devices provided an extremely challenging problem from the numerical point of view, and a similar observation can be made for the area of electronic circuit simulation in the 80s and 90s. Owing to the advent of new numerical methods, these simulations are now carried out routinely.

A relatively new development within the electronics industry is the use of software for the simulation of interconnects structures connecting the different components in a large integrated circuit. Figure 1 shows such a structure. In present-day technologies, this structure may contain up to seven layers.

With shrinking dimensions and increasing frequencies, nearby metal tracks and layers can have an influence on each other. So far, these effects could

be neglected when analysing the behaviour of an integrated circuit. However, there is clear evidence that this will not hold for future generations. Electromagnetic effects (see Figure 2) may influence the functioning of a circuit to the extent that signals are delayed intolerably. Clearly, this must be avoided, explaining the large interest in the topic of signal integrity.

In close cooperation with Philips Research in Eindhoven, the Scientific Computing Group at the TU Eindhoven started an investigation into numerical techniques for the simulation of the problem. A European project (CODESTAR) was set up with two challenging tasks. Firstly, to develop efficient numerical techniques for solving the three-dimensional electromagnetic problem. Secondly, based on these calculations, to construct models with a small number of parameters that adequately describe the electromagnetic




Figure 1: Interconnect structure.

behaviour. This step is commonly referred to as 'reduced order modelling'. It has become clear that methods available for these tasks need a thorough revision in order to be able to cope with present demands of the electronics industry, and so the Scientific Computing Group concentrates on both aspects.

The time-dependent three-dimensional electromagnetic problem is commonly solved using the finite difference time domain method (FDTD), developed by




Figure 2: Electromagnetic effects.

Yee in 1966. The method is simple, and programming it is rather straightforward. However, there are some clear disadvantages. First of all, the time stepping is explicit, so the smallest spatial dimension dictates the time step. In addition, the spatial grid is often unnecessarily fine due to the finite difference character. The Scientific Computing Group has developed a general framework of time stepping methods, which relieve the strict stability conditions. This is extremely important for future simulations of signal integrity. Also, we are now working on other improvements, such as the use of parallel computing, combining two solutions on different staggered grids to improve the accuracy, and the introduction of local mesh refinement using local defect correction (LDC).

The second main task is even more important than the first. In fact, reduced order modelling is a discipline, which receives growing attention in many application areas as well as in mathematics. To stress the importance, the Scientific Computing Group has organised several international workshops within the framework of the European network MACSI-net. Reduced order modelling aims at capturing complex behaviour into a compact model with a small number of parameters. The following sequence of figures clearly illustrates the intention: although only a few straight lines and planes describe the bunny on the right, it is clearly recognized as being a bunny.




Figure 3: Principle of reduced order modelling.

In the past year, we have developed a robust reduced order modelling technique based on an expansion in Laguerre functions, combined with Krylov subspace techniques. The method leads to low order models that are stable both in the frequency domain and the time domain, and has been implemented and tested in a Philips Research software tool for the simulation of electromagnetic compatibility. In the next year, it will be used within the framework of the CODESTAR project to generate compact representations of the electromagnetic behaviour of interconnects.

Links:

- <http://www.win.tue.nl/scg>
- <http://www.imec.be/codestar>
- <http://www.macsinet.org>

Please contact:

Wil Schilders, Eindhoven University of Technology, The Netherlands
 Tel: +31 40 247 8389
 E-mail: w.h.a.schilders@tue.nl

Contpack — Maximising the Volume Utilisation of Containers

by Onno Garms, Ralf Heckmann and Stefan Rank

The Fraunhofer Institute for Algorithms and Scientific Computing (SCAI) has developed three-dimensional packing software for optimising the volume utilisation of containers. The software automatically chooses from a given range of containers and quickly generates highly efficient three-dimensional packings of complex objects.

Today's industrial production is often spread over many sites and locations worldwide. Manufacturing of various parts and intermediate or final assembly can all occur in different places, and the integration of these production sites requires efficient transportation of the turnout between them. This transportation of parts or products is often done in

containers or boxes. In order to minimise the cost of transport and logistics, it is beneficial to many companies to pack their products as tightly as possible.

The optimisation group within Fraunhofer SCAI is currently developing automatic packing algorithms and software that can be used to pack complex

three-dimensional parts (each of which may have many hundreds of thousands of facets) into containers as efficiently as possible, thereby maximising their volume utilisation.

Apart from volume utilisation, the results calculated by our packing algorithms are subject to other constraints

Courtesy by the online magazine 'logistics pilot'.



Figure 1:
Workers loading containers with parts.

handle CAD data of different quality. The integration of the Contpack program into existing CAD systems allows both the design of new components and the planning of their transportation to be performed in a single step within a uniform environment.

The user can select a triangulated VRML representation of the component to be packed and a list of container sizes to be tested, and can define the stable positions. The stable positions restrict the degrees of freedom for the components to rotations around the vertical axis and to arbitrary translations. This accounts for the stability of the generated layout.

There are a number of packing options, including lattice packing. The lattice is a two-dimensional grid made up of repeating units such as rectangles, trapezoids or hexagons. The lattice maintains the layout during transportation and thus protects the components. The use of a lattice results in a completely regular packing of the container. For achieving denser layouts there are options without a given lattice structure. The resulting layout can be viewed in three dimensions, allowing the user to browse through the tiers of the filled container.

Packing tasks belong to the class of NP-hard problems, meaning that no algo-

and demands from industry. Since the loading and unloading of containers is often done by humans, it is important to generate layouts which are easily reproducible by humans. The number of different rotational positions of the components is therefore limited to a small number. These rotations are not provided in advance, but must be found and optimised by the software. Additionally, the containers are filled in tiers, many of which are identical.

In current industrial practice, the generation of packing layouts for containers is done by human beings, who require a sufficient number of physically available components and containers, and complete the task by trial and error. Such a procedure is impossible if the components have not yet been produced. Increasingly however, companies would like to generate container packings for their parts while still in the design stage, in case there exists a way of modifying the outer hull of a part slightly such that it remains equivalent from a technological point of view but can be packed more efficiently.

Finding optimal packing methods with physical objects is a difficult and time-consuming task, because both component geometry and a large variety of possible container dimensions must be considered. Another common method uses CAD systems to generate the layout, but these lack algorithmic support for the task. As things stand, the packing quality of a part into a container improves over time as the insights and suggestions of employees are implemented.

Our Contpack software chooses a proper container size from a given list of sizes and generates a high-density regular layout for this container. This reduces the overall number of containers needed in the supply chain, including stocking, transportation and handling. In the stage of production planning, the automatic packing generation helps estimate the required number of containers. It can also help in recognising the need to develop containers with new dimensions.

In contrast to existing algorithms, which are often limited to the packing of cuboids using bin-packing strategies, our algorithms can handle components of arbitrary shapes. The software can be used at an early stage in component development, because the algorithms can

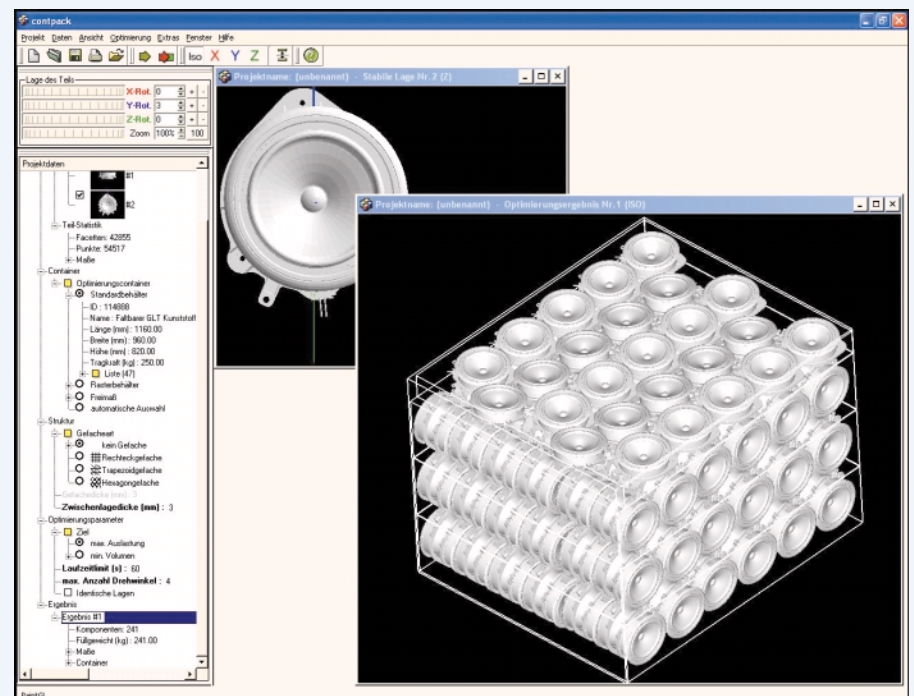


Figure 2: Computer-generated packing of loudspeakers into a container.

rithm for the optimal, time-efficient solution of the problem is known. The techniques employed include exact and heuristic optimisation methods depending on the complexity of the task or sub-task being performed by the algorithm. For packings into given lattice structures we have implemented very fast optimal methods. For packings without lattice structures, the problem is NP-hard. We also apply strategies for finding regular packings, but we add heuristic methods for finding optimal combinations of regular packings in different areas of the box.

Future work will include relaxing the strict separation of the tiers inside the container by robust intermediate tiers. Also, more than one stable position per tier will be allowed. The user will be able

to define taboo areas between the components, and we will fully integrate the software into the user's CAD environment, eg into CATIA.

Fraunhofer SCAI has a strong background in the solution of cutting and packing problems. Experts in the department for optimisation have been working on packing problems for over ten years. The application of the software to three-dimensional container packing started in 2002.

Fraunhofer SCAI is collaborating on this project together with and by the commission of Audi and the BMW Group. Our partner Solve-IT MVI in Munich took over the project management and developed the interface to the container and VRML data, the dialogues, the visualisa-

tion in 3D and the generation of the packing instruction reports. While the program is tailor-made for the needs of the automotive industry, it can be used in many areas of the manufacturing industry. Fraunhofer SCAI intends to release a standard product for the container-packing problem in the near future. We also plan to extend the functionality of the software and to make it usable in other fields of application.

Link:
<http://www.scai.fraunhofer.de/opt>

Please contact:
 Stefan Rank, Fraunhofer SCAI
 Tel: +49 2241 14 1503
 E-mail: Stefan.Rank@scai.fraunhofer.de

First Commercial Bayesian Software for Intelligent Troubleshooting and Diagnostics

by Claus Skaanning

Dezide Advisor is the first system of its kind based on Bayesian networks – a modern approach for artificial intelligence, providing better solutions for complex and uncertain problem domains. The system is able to provide intelligent advice for solving complex technical problems efficiently and precisely, and no special expertise or skills are required to develop the underlying knowledge base.

Dezide's product, Dezide Advisor, implements state-of-the-art research within the area of intelligent troubleshooting and diagnostics, and is well suited for support and maintenance of complex industrial machinery.

Dezide is a small Danish company staffed by former employees from an R&D department of Hewlett-Packard that was 'spun off' in 2001. There are currently eight employees, including one doctorate and three Masters, all specialists in computer science with a focus on Bayesian networks.

Bayesian networks allow the structuring of complex information in a graphical network, where arrows depict the direction of causality, that is, from cause to effect. The exact computation methods

of Bayesian networks then allow the user to fix any variable, and continually recalculate the probabilities of other variables given the new observation. In comparison with fuzzy systems and other AI technologies, Bayesian networks are mathematically sound, allowing them to be used for far more complex and uncertain problems.

Dezide's software is currently being used for troubleshooting in many different areas, including Internet connection problems and email setup problems (TDC, Denmark's largest telco), network router problems (Telenor), and common IT problems (SAS). In addition to this, Dezide is currently involved in activities where the software is being used for the maintenance of complex bridge structures and diagnostics of windmill systems.

Dezide Advisor is based on Bayesian networks and implements seven patents in this field. In short, a field expert stores his knowledge of problems, resolutions and symptoms using a special authoring tool. This knowledge can then benefit end-users, technicians, and operators in trying to more quickly and precisely resolve problems. In addition, Dezide Advisor can help with detecting deviations from normal operating behaviour, can predict future difficulties based on current measurements, and can find the probable cause of a problem and then either sort it out itself or call for the necessary assistance.

Dezide Advisor is particularly well suited for troubleshooting and diagnosis in complex technical systems. In this case, Dezide Advisor will guide the user

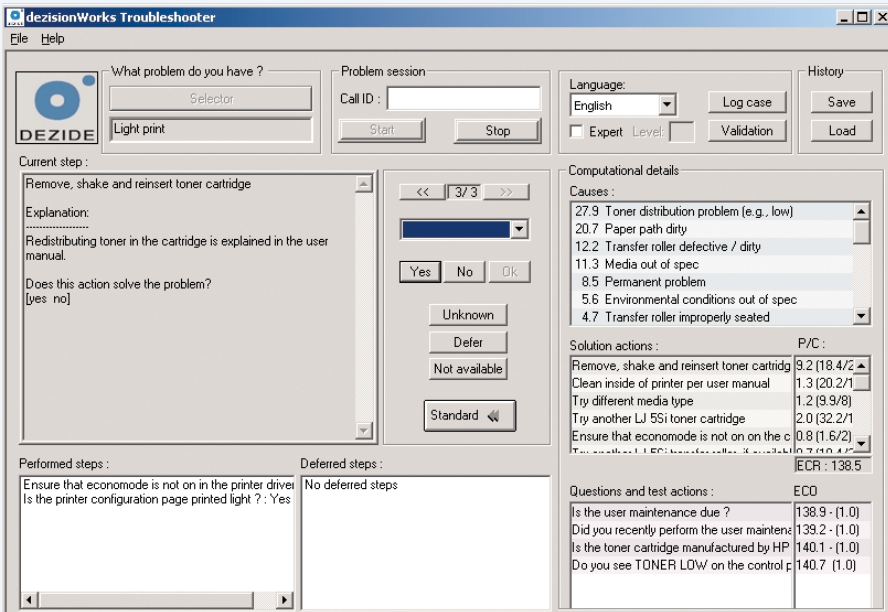


Figure 1: Troubleshooting printer problems.

towards solving a problem or finding a diagnosis

- validation tools for testing and quality assurance.

Dezide Advisor includes a diagnostic engine that can be embedded in diagnostics applications, desktop tools for constructing and maintaining the knowledge base, and a scalable, secure and robust server system for Web-based access by thousands of simultaneous users.

It is possible to learn to use the knowledge engineering tool in only two days, and no knowledge of Bayesian networks is required. The tool includes a graphical visualisation functionality to get an overview of the knowledge base (see Figure 2).

through an optimal sequence of steps leading to a solution or diagnosis in as few steps as possible. Dezide Advisor can run interactive question-answer sessions, but can also run fully automatically, embedded in technical systems.

The next step in an interactive question sequence is calculated based on all previous answers, probabilities that steps will help, and their costs (time, risk, money etc).

Dezide Advisor is an improvement over existing solutions in several ways. It is able to reach a solution or diagnosis in as few steps as possible and with the lowest possible cost, and when it concludes, it

presents a list of the potential root causes sorted by probability. The system is very flexible, as the user is not required to follow a fixed path but can skip steps that he cannot answer, and jump directly to steps that he wants to answer straight away. In addition to this, the Advisor is self-adapting and adjusts itself to the situation over time.

Dezide Advisor is based on several breakthrough inventions. These include:

- an authoring tool that allows anyone to construct Bayesian networks efficiently for diagnostics and troubleshooting
- algorithms for calculating the optimal sequence of actions and questions

Dezide is now looking at improving its software to handle more complex problems in industrial plants. Together with ABB, SICS, Aalborg University, the University of Stuttgart, Renault and several other research and industrial partners, Dezide is a member of a consortium proposing an EU research project known as 'The Butler Concept' (see information box on page 16).

In the Butler project, Dezide will be involved in conducting general research for improved knowledge acquisition and troubleshooting algorithms within highly complex industrial plants. In particular, Dezide will work on methods for tuning the knowledge base based on expert feedback, generating easily understandable explanations of the system's suggestions and conclusions, and importing of diagnostic data from other formats such as log files, decision trees, case bases and so on.

Link:
<http://www.dezide.com>

Please contact:
 Claus Skaanning, Dezide, Denmark
 Tel: +45 9655 7370
 E-mail: claus.skaanning@dezide.com

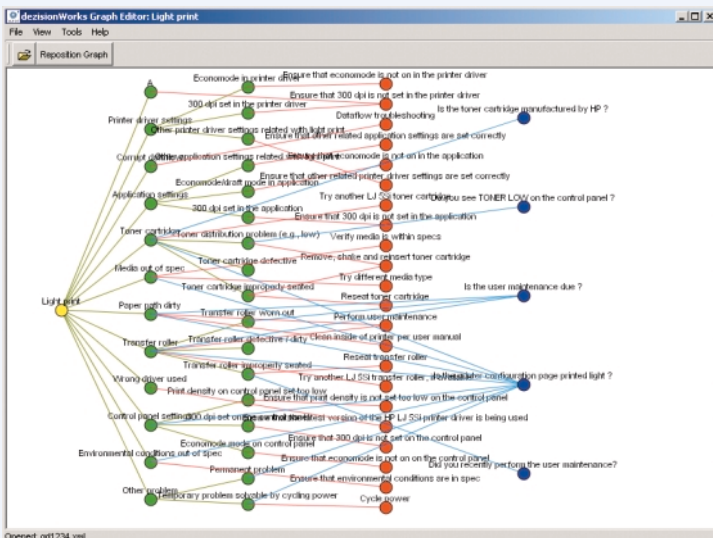


Figure 2: A graphical view of an example knowledge base.

Innovating Shoe Manufacturing using Advanced Simulation Techniques

by Emanuele Carpanzano and Andrea Cataldo

In order to handle the increasing complexity of distributed industrial manufacturing systems, there is a strong demand for methods and tools that support the designer in the analysis and optimisation of flexible infrastructures for the automation of production processes. We are currently investigating how emerging standards and advanced simulation techniques can be exploited successfully in the production of custom made shoes.

Nowadays, many different approaches, methodologies and computer aided design tools can be employed in the analysis and optimization of complex industrial systems. As a consequence, the selection of the most appropriate techniques and tools for a particular system is critical. Before making any final decision, it is necessary to investigate the functionalities offered, compliancy with existing standards, and capacity for interoperability with non-proprietary instruments. We are currently working on a methodological approach to the structured design and simulation-based analysis and optimisation of manufacturing plants. The methods and tools used during the different phases of the design are illustrated with reference to an application in the shoe manufacturing sector. In particular, we describe an innovative plant for the production of customized shoes. This plant is located at the ITIA-CNR laboratory in Vigevano (Italy) and constitutes the pilot plant of a large European research project called EUROShoe, which began in March 2001 and is scheduled to finish in June 2004.

EUROShoe is a complex and ambitious project with thirty-five academic and industrial partners from ten different European countries. It aims at a dramatic renovation of the concept of shoes as products and of their production, based on a transformation from mass-produced to mass-customised goods. This product evolution goes in parallel with a transformation of footwear companies into distributed and flexible enterprises capable of handling the complexity that such a radical change in the nature of the product implies and of mastering the associated new technological challenges. This implies a complete revision not

only of the entire manufacturing process but also of the tools used to analyse and optimize the resulting innovative industrial system.

For the sake of brevity, the methods and tools used for the design, analysis and



Figure 1: The Innovative Shoe Manufacturing Plant of the ITIA-CNR laboratory of Vigevano (Italy).

optimization of the plant automation system are outlined with reference to the control and supervision of the transport line. In the shoe manufacturing plant that we consider (Figure 1), an innovative transport line is used to move the semi-finished shoes from one machining station to another according to a predefined operating schedule. In particular, the innovative molecular structure of the transport line (Figure 2) strongly enhances the modularity, scalability, integrability and reconfigurability of the production system, thus increasing the overall flexibility of the plant.

At the beginning of the design life cycle, the control and supervision system of the transport line was specified by means of UML (Unified Modelling Language) diagrams; in particular use-case, class, sequence and state diagrams were used to specify the systems modules, their relationships and their dynamic behaviour. The system architecture and functions were designed using the Function Block formalism defined by the IEC 61499 standard, which integrates object-oriented concepts and discrete event models to suitably support control software design. The functional model obtained was analysed and optimized by means of closed-loop discrete event simulations performed in the Simulink/Stateflow environment, where both the controller and the controlled process were modeled and simulated through state charts, as was the behaviour of the overall system. In this way, the correctness and performance of the solution proposed could be evaluated. In particular, a bottom up methodology was defined and exploited to study the system according to a modular approach. This enabled us to simplify the overall analysis process, to verify the correctness of the automation functionalities easily and to optimize the system performance from the very first phases of the design life cycle. Once the functional modules had been verified, the control and supervision algorithms were developed using the SFC (Sequential Functional Chart) formalism, which is an advanced discrete event modelling graphical language, directly derived from Petri nets, and included in the IEC 61131 part 3 standard, which defines programming languages for industrial PLCs (Programmable Logic Controllers). The SFC algorithms were

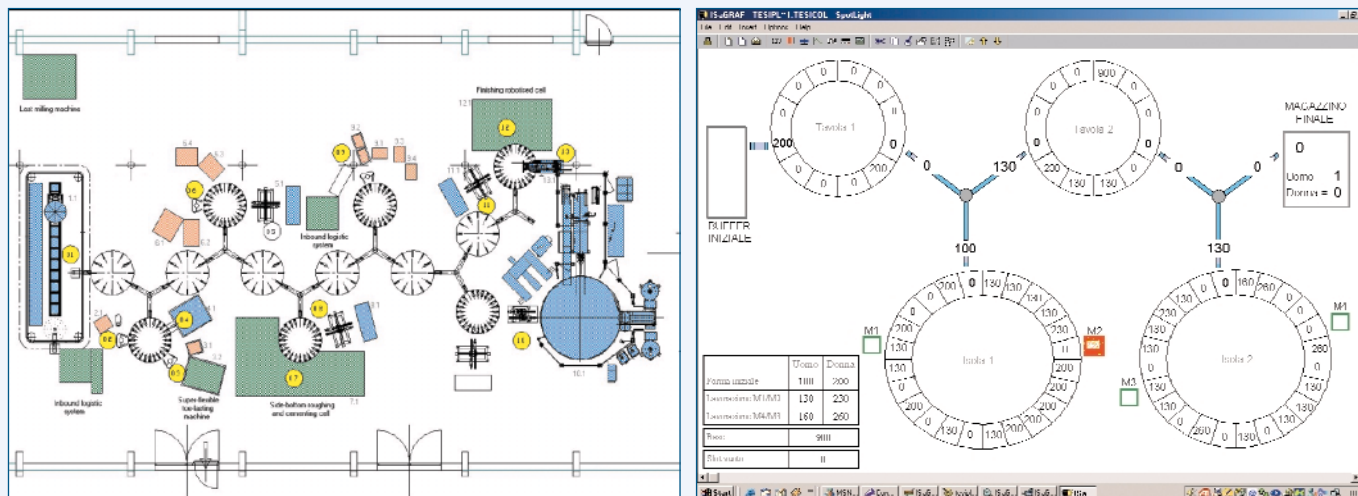


Figure 2: Layout of the Shoe Plant (left) and ISaGRAF Simulation Graphical Interface (right).

analysed through closed loop discrete simulations in the ISaGRAF environment in order to verify their correctness and to optimize their performance. ISaGRAF is a CACSD (Computer Aided Control System Design) tool that allows the automation software to be structured using all five IEC 61131 languages and supports simulation functionalities for testing purposes, as well as automatic code generation facilities for different industrial systems. In order to perform the closed loop simulations, simplified models of the plant devices have been represented in ISaGRAF by means of suitable SFC modules and data struc-

tures. Moreover, to simplify the analysis of the simulation results, a simple 2D graphic animation was realized in ISaGRAF (Figure 2). Simulations have been performed by considering typical operating conditions, ie, typical production orders, and the results obtained show that the system is deadlock free and that the plant is well balanced, ie its resources are all used effectively. After verification and optimization of the automation functions and algorithms, the corresponding code was generated and implemented on the target industrial devices. We found that the simulation-based analysis techniques reduced the

plant rump-up times and costs and improved overall system performance. Future work will concern the exploitation of the standards, methods and tools proposed for other industrial plants and their integration within an advanced CACSD tool for manufacturing systems.

Links:

<http://www.itia.cnr.it/>
<http://www.euro-shoe.net/>

Please contact:

Emanuele Carpanzano, ITIA-CNR
 E-mail: e.carpanzano@itia.cnr.it

Billy Goat Detects Worms and Viruses

by James Riordan and Diego Zamboni

Researchers at the IBM Zurich Research Laboratory have designed and implemented a new intrusion detection tool that not only facilitates early detection of attacks but also sharply reduces the false alarm rate.

The Billy Goat is specialized to address the problems posed by network service worms. As an intrusion detection sensor, its most important property is that it is free from the high rate of false alarms produced by many other sensors. It achieves this property through the use of a novel architecture that combines an extensive view of the network, spoofed service interaction with potential attackers, and a clear focus on detecting automated attacks.

One of the greatest threats to security has come from automatic, self-propagating attacks such as viruses and worms. These attacks scan at random until they are able to place a program on the server using a maliciously crafted request. The program uses the now-infected server as a base from which to attack other servers. The direct result is rapid exponential growth in the number of attackers leading to load-induced network failure.

While the presence of these attacks is by no means new, the damage that they are able to inflict and the speed with which they are able to propagate has become paramount. Further increases in connectivity and service complexity only threaten to exacerbate their virulence.

Architecture

The Billy Goat functions by spoofing the existence of machines and services at otherwise unused IP addresses. Because the addresses are otherwise unused, all

traffic destined to them is a priori suspicious. The sensor spoofs services, rather than merely recording attempted connections, to determine the intention behind the traffic.

Billy Goat is built atop a security-hardened Linux machine that offers no real services beyond very restricted login. It is configured in conjunction with the network on which it runs so that traffic directed toward address subnets that are not used is routed to the Billy Goat.

The Billy Goat itself offers a virtualization infrastructure that allows individual sensors to be written as if they were running on a single host. It also provides a logging infrastructure based on a relational database facilitating correlation and analysis of the copious data produced by the large number of virtual sensors.

The spoofed services include HTTP, Microsoft RPC (remote procedure call), Microsoft SQL (database), and SMB (file sharing and printing). Vulnerabilities in these services are commonly used as vectors for worm and virus propagation.

One very important requirement of the Billy Goat is that it continue to function in times of heavy worm activity. In particular, it must retain its utility even when the performance of the network is dramatically diminished or even completely unreliable. This requirement implies a distributed system of Billy Goats on the network. Each local Billy Goat serves to inform administrators of local infection so that it can be eliminated. In order to support the distributed architecture, the Billy Goat distribution contains an automatic update mecha-

nism. This mechanism ensures that a deployed Billy Goat has all the latest sensors and signatures.

It is worth contrasting the Billy Goat Box with another security tool called a Honey Pot. A Honey Pot, as its name suggests, offers "something desirable" to lure attacks to a particular machine. Unlike a Honey Pot, the Billy Goat does not advertise any services, and does not allow any of its services to be corrupted. The Billy Goat is designed to detect large-scale automated attacks rather than trying to deceive sentient human attackers.

Please contact:

James Riordan, Diego Zamboni
 IBM Zurich Research Lab/SARIT
 Tel.: +41 1 724 89 81
 E-mail: rij@zurich.ibm.com

Decision Support for Personnel and Resources Management in Multi-Project Environments

by Antonia Bertolino, Eda Marchetti and Raffaella Mirandola

Decision making in project management requires accurate and automated support to monitor and control the development process. Managers must make reliable schedule predictions and optimize personnel utilization. They thus must be able to dynamically evaluate whether the resources assigned to a job are sufficient and whether the organization structure is adequate to meet the scheduled deadlines. To support managers in these tasks, scientists at ISTI-CNR are studying the potential of classical performance engineering and propose the Propean (Project Performance Analysis) methodology.

Propean (Figure 1) relies on the metaphor that project teams correspond to the processing resources of classical performance models, and project activities are the tasks to be performed within established time intervals.

According to this metaphor, well known techniques from performance analysis can be usefully adapted to managers' needs, for eg assessing the time for completion of specified tasks, handling personnel multitasking over concurrent projects, distributing the workloads in development cycles, deciding about product release, and similar issues.

However, a key aspect for the effective deployment of Propean is to adopt an interface that appears familiar to managers. Hence, we adopt as the input modelling notation the UML (Unified Modelling Language) Profile for Schedulability, Performance and Time (RT-UML) recently approved as an OMG standard. In this way, managers do not need to become knowledgeable of performance notations, but can develop a model of the flow of activities and of tasks distribution using the widespread UML notations and tools. The translation into a format that can be processed by performance analysis algorithms is then performed automatically.

The SPE Approach

The Propean technology is based on the Software Performance Engineering (SPE) approach, which is a systematic, quantitative approach to software system construction in order to meet performance objectives. Its basic concept is the separation of the Software Model (SM) from its environment or Machinery Model (MM).

The SM captures the essential aspects of software behavior; it can be represented by means of Execution Graphs, in which nodes are software workload components and edges represent transfers of control. Each node is weighted by use of a demand vector that represents the

resource usage of the node. The MM models the hardware platform and is based on Queueing Network (QN) models.

The Methodology

Starting from the above metaphor we capture the aspects relative to the activity planning in the SM, and those relative to people utilization and distribution in the MM. These models are derived from Sequence Diagrams (SD) and Deployment Diagrams (DD) respectively, and are automatically combined by Propean which generates a unique QN, that can be analysed by traditional performance analysis tools.

The methodology is currently being implemented. Figure 2 illustrates the tool architecture including: a UML case tool (specifically ARGO) which processes the input diagrams and generates an XML file; a component which takes as input the XML file and produces the corresponding QN; a QN analytical solver or simulator, which derives the performance analysis results; and finally, a results converter which analyses the performance results and converts them back into UML.

Here below, we outline the basic steps of the Propean methodology:

1. Analysis: definition of project activities (manager)
2. Modelling: definition of the SDs and the DD (manager)
3. Model annotation: specification of parameters and values (manager)
4. SPE models generation: derivation of the SM and MM models (automatic)




Figure 1: Different approaches to decision-making.

5. Model evaluation: resolution of the QN and derivation of the relevant predictions (automatic)
6. Analysis of results: evaluation of the results obtained (manager).

Applications

One of the advantages of using performance techniques for project management is their natural capability to handle multiple projects and their mutual involvement in schedule and resources usage. In fact, Propean results particularly useful in environments where multiple projects compete for common resources, which is for instance the case in most software development environments.

Propean is the result of a collaboration between ISTI-CNR in Pisa, the University of Rome Tor Vergata and Ericsson Lab Italy (ERI) in Rome. It has been applied with encouraging results to some industrial case studies. Here we mention two of them: one is the widespread Rational Unified Process (RUP) customised for a specific ERI project. Considering RUP itself as a product, Propean provided answers to the following questions: How long will RUP take to process a certain project? How will RUP utilize the available resources? How is the RUP schedule affected by the concurrent processing of several projects? The application of Propean to RUP was divided into two steps: the first, called RUP modelling, consisted of the description of the functionality and architecture of the RUP product by means of UML diagrams annotated according to the RT-UML profile. The second step, called RUP customisation, represented the core of Propean application and included the refinement of the UML diagrams according to personnel availability and process requirements.

In a second case study, Propean was used to support the decision to release a software product, after testing is completed with no trouble report left open. Propean can assist managers in the early prediction of the expected time until release, considering the actual resource availability (with personnel multitasking). Alternatively, if the release time is fixed with the customer, Propean can help decide the most adequate personnel configuration to fulfil the time constraints.

Propean can provide useful support to decision makers in project management. The tool can provide valuable input for analysis, but human expertise remains essential to model the processes, to tune them with the correct values and to interpret the results obtained.

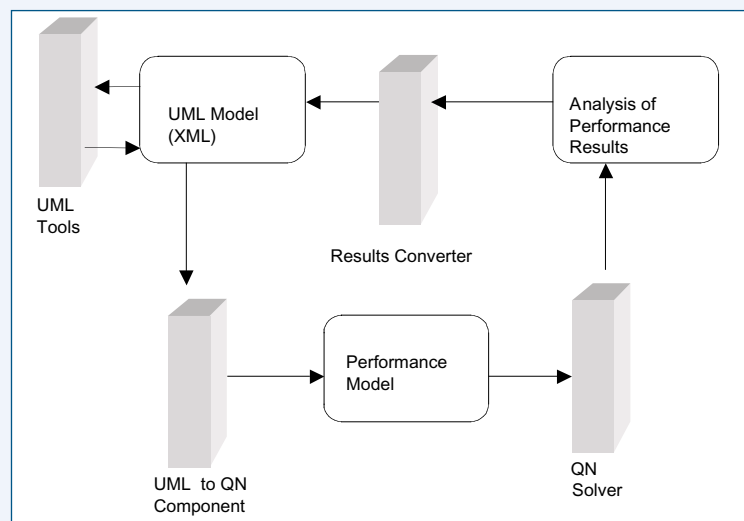


Figure 2: Propean Tool Architecture.

Links:
<http://www.isti.cnr.it/ResearchUnits/Labs/se-lab/index.html>
<http://www.omg.org/technology/documents/formal/schedulability.htm>
<http://www.perfeng.com/>

Please contact:
 Eda Marchetti, ISTI-CNR
 Tel: +39 050 315 3467
 E-mail: eda.marchetti@isti.cnr.it

Towards Useful Tools for Estimating Railway Infrastructure Capacity

by Jan Ekman

With the aim of supporting the process of adapting the railway infrastructure to present and future traffic needs, SICS is developing methods and tools for estimating railway capacity. The project is supported by Banverket, the Swedish National Rail Administration.

Traffic on the Swedish railway network has increased considerably in recent years, and large investments in railway infrastructure are planned in order to meet the expected need for higher railway capacity. When new infrastructure is designed, it is important for the quality of the transportation system that the infrastructure and the traffic solutions are adapted to each other. For this purpose, new standards, methods and tools are required for precise analysis of infrastructure capacity.

Our starting point is that railway traffic is to be regarded as a set of traffic patterns that are repeated on sections of the track system. From this standpoint we have developed a method for analysing capacity in signalling design. Consider the left side of Figure 1, which depicts an instance of periodic traffic, by which we will describe the method. Each period of this repeated traffic pattern consists of three trains (w , e and f) traversing a single line with three meeting stations ($ms1$, $ms2$ and $ms3$). Each of the advances of the two trains w and e are separated into two parts, $w1$, $w2$ and $e1$, $e2$ respectively, by the waiting points at

which the trains wait for other trains to pass. We use the term 'cleared paths' to indicate the paths of the trains produced by this separation. Waiting points and cleared paths are the basic concepts on which the method relies.

The 'traffic pattern graph' on the right in Figure 1, shows the priorities of the movements on the cleared paths in the traffic pattern. There is one node in the graph, for each 'cleared path'. Straight arcs concern priorities in one and the same period, and bent arcs concern priorities in two consecutive periods. For instance, the bent arc from $e2$ to $w1$ means that the movement $e2$ in period n precedes the movement $w1$ in period $n+1$. The arcs only concern direct priorities.

The compact model of the traffic pattern given by the traffic pattern graph is the first step in the construction of a set of conditions to be used for estimating capacity. The next step in this construction is the addition of arcs for movements that may directly impede on each other. This is followed by the decoration of each arc with the shortest length of

time between the starting times of the two movements. The result, for some invented shortest durations between starting times, is the 'conflict graph' shown in Figure 2.

Each arc in the conflict graph is a condition. For instance, the bent arc of weight 25 from f to $e1$ represents the condition that the movement e in period $n+1$ may not start earlier than 25 time units after the start of the movement f in period n . Given the conflict graph, the minimum period is obtained by analysing the circuits in the conflict graph, where a circuit consists of one or more arcs. The cycle mean of a circuit is the sum of the weights divided by the number of bent arcs, and the minimum period of the traffic pattern equals the maximum of the cycle mean for all circuits in the conflict graph. The maximum cycle mean in a conflict graph can be efficiently computed, with running time $O(n^3)$, at worst, with n as the number of cleared paths. The minimum period for the considered traffic pattern is 50, where the circuit of weight 100 consisting of the edges $(w1,w2)$, $(w2,e1)$, $(e1,e2)$ and $(e2,w1)$, two of




Figure 1: Traffic pattern and traffic pattern graph.

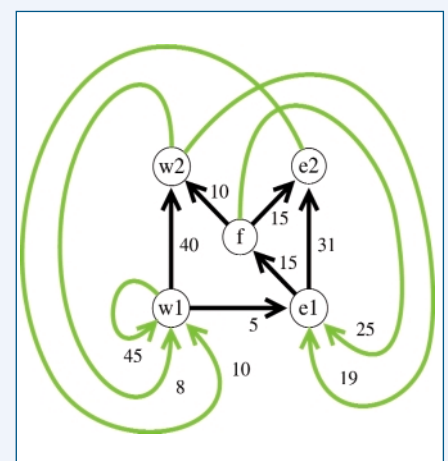


Figure 2: Conflict graph.

which are bent, is a critical circuit. That is, a critical circuit is a circuit with cycle mean equal to the maximum cycle mean.

To sum up, this method allows us to decide the minimum period for a given traffic pattern on a railway section. This is in itself useful for deciding on the best infrastructure for the planned traffic, given, for instance, by a timetable. The method applies to precise examination of infrastructure, such as considering the effects of adding and altering the positions of track signals. The method suggests that capacity of a railway

section can be estimated as the overall minimum period for a collection of automatically generated traffic patterns, each of which satisfies some requirements on the traffic. It may also be of interest to know, for a given fixed infrastructure, which traffic solutions are suitable and which are not. In such a case, rather than simply estimating the capacity, it is also necessary to determine which traffic patterns are badly adapted to the infrastructure, ie, those traffic patterns with a long minimum period. Without paying any attention to robustness of the traffic system, analysis of capacity would be

merely of theoretical interest. One way of dealing with robustness is to add default buffer times to the course of events on which the model is built, making this set of buffer times the robustness criterion. However, the model of the traffic pattern given by the conflict graph is well suited as the basis for robustness analysis, as for instance stability analysis and delay propagation.

Please contact:

Jan Ekman, SICS
Tel: +46 8 633 15 65
E-mail: jan@sics.se

Efficient Utilisation of Resources in the Railway Industry by Levelling Departures and Arrivals of Trains

by Martin Aronsson

The raison d'être of the transportation industry – moving things between locations – is also the source of a key problem; namely, that the resources necessary for operation are geographically distributed. This means scheduling can become an extremely complex problem, particularly in comparison to stationary industries, specifically traditional job-shop scheduling.

In the railway industry, the traditional method has been to first produce the timetable, and then to assign resources such as engines, cars and personnel to the scheduled transports. Nowadays assignments are done with optimisation tools that are based, for example, on network flow models.

However, since the deregulation of the railways in Sweden, and due to the limited capacity of the track resources, competition has increased among operators for attractive timetable slots. Although the commission of the European Union has stated that the European railway shall be deregulated, different countries have reached different developmental stages. Sweden has deregulated cargo traffic and some passenger traffic, and has separated railway management and slot allocation from traffic operations.

The competition for track slots comes partly from varying customer demands



Photo: Green Cargo.

Figure 1: Efficient resource utilisation is a key factor in the railway industry.

on the railway operators, and partly from the growing interest in the efficient usage of expensive resources such as engines (see Figure 2). It is increasingly apparent that there are two ways to efficiently use resources, namely, optimal assignment of resources to tasks, and well-scheduled departure and arrival times, which facilitate efficient assignment. Before the application for timetable slots is made, each operator

must find the best compromise between resource usage and customer demands. This stage in the planning process, made critical by deregulation, currently suffers from a lack of good decision-support tools, and improvements are necessary in order to facilitate good assignment in later stages.

Over the last few years, SICS has addressed this early process stage. For a human, the variables of geography, time and assignment mean the problem rapidly becomes very complex. To help practitioners, SICS has developed a series of computational models for this process, in particular one constraint-programming model and one integer-programming model. The latter is currently being integrated into Sweden's largest cargo operator. The model is an extension of a network flow model, in which time windows have been incorporated, leading to a mixed-integer program. By relaxing the departure and arrival times to be within lower and

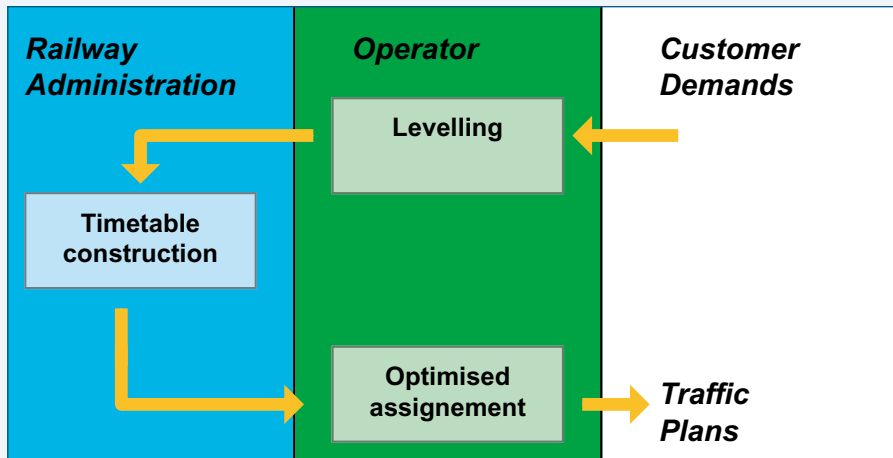


Figure 2: Process flow, customer demands and resource utilisation requirements restrict the timetable process.

upper bounds, and adding constraints that ensure that certain relations between trains still hold (eg for cars that should be decoupled and attached to another train), the tool can produce efficient rotations for the engines while moving the transport operations inside the time window. In the general case, a network flow model with time windows becomes very complex, but in practice it turns out that we can compute solutions to the generalised problem as long as the time windows are not too large. Using itera-

tive application of the process, we have successfully computed solutions for time windows up to sixty minutes.

In the stationary industry, this process of moving tasks in time to reduce the resources needed is sometimes referred to as 'levelling', since it tries to level the resource usage over time. While this is a simpler task in stationary industries, it is just as important in the transportation industry. In Figure 3, a typical resource graph is presented for a railway cargo




Figure 3: Resource utilisation as a function of time (5 days) with peaks every night.

company. It is the peaks that set the number of resources (engines) needed for the transportation stock. By using the levelling software, the resource peaks can be considerably reduced. Our full-scale tests have shown that 5-10% resource savings are within reach.

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

Please contact:
 Martin Aronsson, SICS
 Tel: +46 8 63315 87
 E-mail: martin@sics.se

Real-Time Optimisation of Empty Rail Freight Car Distribution

by Martin Joborn

The rail freight carrier Green Cargo has a new system for rail freight car distribution planning. It is an automatic, real-time planning system, based on a kernel of optimisation models and methods. The new system has resulted in both significant savings and quality improvements.

In railway systems, it often occurs that there is an imbalance in the freight flows, meaning empty freight cars must be moved in order to enable new transports. An efficient empty freight car distribution process is vital for good utilisation of the freight cars. A bad distribution process can result in delayed transports, good-will losses and an unnecessarily large car fleet with the associated capital and maintenance costs.

Green Cargo is the largest railway cargo carrier in Sweden. In collaboration with Carmen Consulting and Linköping University, Green Cargo has developed an advanced new system for planning their empty freight car distribution. The aim of the system is to improve the quality of the empty freight car distribution process. The system handles about 7000 freight cars, divided into 25 types. The railway network has about 1500 pick-up and delivery points for freight cars, and the system produces detailed

transportation plans between these points.

The kernel of the new system is an optimisation model that calculates distribution plans in real time. The model is constantly updated and re-optimised as new information becomes available, eg updated freight car status, transport orders or train capacities. The response time of the system for re-optimising the plans is less than one second. The plans always reflect the optimal distribution

for the current situation. The objective of the distribution planning is to maximise the number of filled transport orders while minimising transportation costs, as well as to generate distribution plans with a preferred structure and to ensure that cars arrive punctually to customers.

The high performance of the system makes it possible to respond immediately to customers' transportation demands. While on the phone placing the transport order, the customer will also be informed of whether freight cars are available for the transport. The distribution system generates an online optimisation to calculate and guarantee the delivery of the empty freight car to the transport. However, regarding the method of satisfying a customer's freight car demand, the system maintains flexibility for as long as possible, while simultaneously ensuring that the order will be filled.

The system is fully automatic and needs no manual interaction to execute the calculated plans. At the correct time, stations, terminals and operating crews

receive instructions relevant to their cars and transport orders - bookings, shunting orders etc – from the system.

The introduction of the new system has made empty freight car distribution a trustworthy and highly precise process, which has led to loyalty from both personnel and customers. The reliability of the process decreases the need for backup solutions such as local freight car



Photo: Peter Lyden

Full control over freight car planning at Green Cargo.

inventories. Furthermore, the freight car controllers have a good overview of the freight car supply-and-demand situation over the system's planning period of one week.

The improved utilisation of freight cars has enabled Green Cargo to reduce its car fleet and the associated capital and maintenance costs. Green Cargo estimates its yearly savings at €7.6 million, excluding the benefits from improved customer relations.

The concept used at Green Cargo is applicable to other problems as well. Interesting potential applications include loaded freight car planning, taxi planning, truck allocation, and container distribution.

Link:

<http://www.carmenconsulting.com/includes/ProjectSheets/RealTimeOptimization.pdf>

Please contact:

Martin Joborn, Carmen Consulting, Sweden
Tel: +46 8 555 79 438
E-mail: martin.joborn@carmenconsulting.com

KUORMA: A Collection of APS-Algorithms for Forest Industry Wood Transport

by Juha Savola, Hannu Rummukainen and Olli Jokinen

The KUORMA system (English acronym LOAD) is a combination of APS (Advanced Planning and Scheduling) algorithms and the system infrastructure services used in the field of forest industry.

These algorithms enable the forest industry to optimise the transport of wood to sawmills and paper and pulp mills using a nationwide calculation area. The entire operation includes about 250 trucks, about 100 points of delivery and about 10 000 wood batches daily. The KUORMA system is part of a more extensive information system, which incorporates optimisation algorithms for wood harvesting, transport planning and routing.

The KUORMA system was developed in a commercial project headed by TietoEnator as a systems integrator. VTT Information Technology was a sub-contractor responsible for the development of transportation algo-

rithms. The customer was Stora Enso Finland. Stora Enso uses about 17.3 Mm³ of domestic wood annually, corresponding to 288 000 truckloads.

Transportation is a remarkable source of cost in the forest industry, especially in Finland where the distances are long and the wood ownership includes private citizens. Not only are wood batches of varying sizes distributed over some 3 000 000 road segments, but the price for standing timber is relatively high. In order therefore to reduce wood procurement costs for competitiveness reasons, one must concentrate on transportation costs. Stora Enso had a system developed

by the coalition in the early 1990s that used greedy algorithms. Its performance and the acceptability of the solutions it provided had reached their limits. Furthermore, improvements in technology had made it possible to use a low-cost Linux workstation cluster with a large amount of essential RAM memory.

The development project started in July 2001, and through October 2002 the system was gradually improved through operational use.

The Techniques Employed

The main components of the KUORMA system are the road network infrastructure, the



Figure 1: A truck unloading at a paper mill.

by a weight limit, which represents a hard constraint for the order in which the batches of a load can be collected. The primary objective in the optimisation model is the minimisation of total driving time for load collection.

Load Assignment Model

The assignment model calculates the best destinations for all loads to be transported. The destinations are factories, sawmills and railway stations (for final transportation by train to the factories). Assignment is taken as input transportation costs based on distances, the daily wood requirements of the destinations, the availability of loads per day and regional and team goals.

Wood requirements are described as amounts for different kinds of wood and penalty costs for deviation from the desired amount. The calculation period is divided into an operational period (a few days) and a tactical period (some weeks). In the operational period, destinations are assigned for every load. In the tactical period, standing wood marked for cutting is optionally reserved for certain destinations because the loads are not yet to be transported.

The results generated by the model show the destinations for loads, the transportation day within an operative period, and the opening day for every stand marked for cutting in the tactical period.

Load Routing Model

The third stage comprises the routing and scheduling of trucks for the immediate future. At this stage, the load destination can no longer be changed, but the choice of which loads to transport in the operational period is

load cluster model, the load assignment model and the routing model.

The key information for the road network infrastructure is the nationally administrated and updated road material, which covers the whole of Finland and contains about 3 000 000 segments of road, including small, private forest roads. The central attributes are road class, x-y geometry and length. The basic road material is organised as an undirected graph structure. Because the set of current objects with a coordinate varies all the time, the calculation of distances had to be performed uniquely for each data set. Nevertheless, the operative calculation effort can be reduced using so-called supermatrices (size 15 000 – 20 000 squared) based on all the distinct points at any given time. Supermatrices are calculated in a Linux cluster during idle hours. When the actual matrices are calculated utilising supermatrices, the hit ratio can exceed 95

%, depending of course on the time between those two calculations. The calculation time is typically a few minutes using a 2.2GHz dual-processor PC with 2Gb RAM. The basic solution given by the road infrastructure is the fastest or shortest distance between any two objects using three different road weight classes.

Load Cluster Model

In the first stage of the KUORMA system, batches of wood in the forest are clustered into loads that fit into trucks. This is done regularly as newly cut stands are entered into the database, and old wood is transported away.

Large batches are easy to divide into truck-loads, but due to the large number of private forest owners in Finland, many batches containing several wood types are quite small, and may only be accessible by low-quality, temporary roads. The road quality is modelled

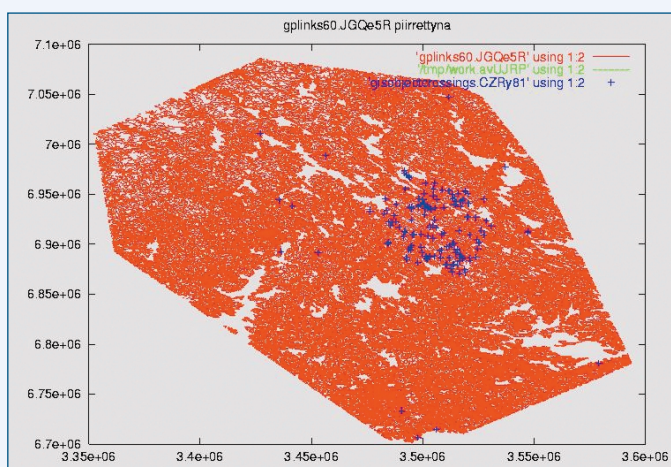


Figure 2: Snapshot from road network of a load cluster model. Before reduction.

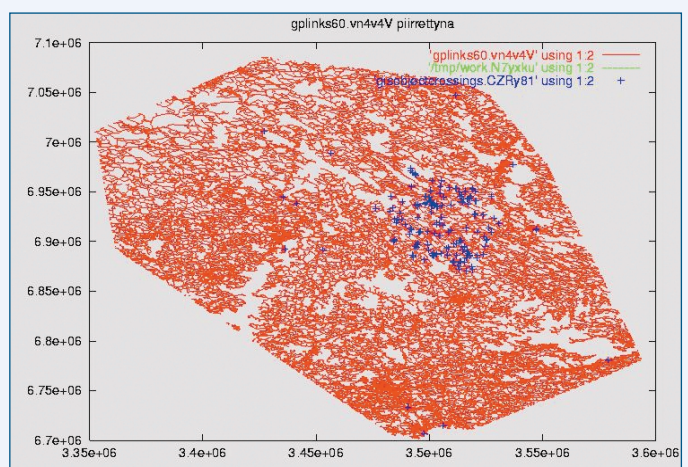


Figure 3: Snapshot from road network of a load cluster model. After reduction.

free subject to certain constraints. The primary objective is to minimise the total driving costs. Trucks may operate all over Finland, but they must regularly return to their home areas to exchange drivers, and may also have scheduled days off. The paper mills and sawmills operate on exacting schedules, and must therefore be provided with specific numbers of loads of different types within given time intervals, although deviations are possible with associated penalty costs. There are a number of other constraints and secondary goals, making this the most complicated optimisation model in the KUORMA system.

Future Activities

The transportation branch of the forest industry stands to benefit greatly from optimisation algorithms. Stora Enso anticipates a 5% saving in wood transportation costs compared to the old, greedy-algorithm approach. Our approach can also be applied to other transportation segments.

Further work on the KUORMA system could include evaluating multimodal transports, adaptation of new road material containing z-coordinates (hills) and making use of daily weather data.

Links:

http://www.vtt.fi/tte/projects/fleet_management/kuorma.html

http://www.vtt.fi/tte/projects/fleet_management/dynatrans.html

Please contact:

Juha Savola, Hannu Rummukainen,
Olli Jokinen, VTT Information Technology
E-mail: juha.savola@vtt.fi

Parc Bandwidth on Demand

by Helmut Simonis

Running demanding services like video conferencing over the Internet is not easy, with packet loss and jitter making it a less than enjoyable experience. The Parc Bandwidth on Demand (BoD) system helps Internet service providers to sell and differentiate high-quality, short-term connectivity for video conferences and other services, thereby avoiding these problems.

The system allows customers to request large amounts of bandwidth for short periods of time, with a firm guarantee of high quality of service (QoS). Requests are processed automatically and, if accepted, are provisioned automatically. The system ensures that requests are only accepted if all QoS guarantees can be met.

A typical example would be a two-hour video conference starting at some fixed time, which requires a reliable (low packet-loss), low latency connection between a number of locations around the world. A requirement might be imposed that the accepted demands must not cause the service level agreements (SLAs) of any existing customer to break. To meet this requirement it is necessary to account for the bandwidth use in the network. The system does this with an automated workflow.

The customer books the BoD service online via a portal up to 24 hours in advance. If the request can be accepted he receives an email confirmation within two hours. At the requested start time, a traffic-engineered (MPLS-TE) tunnel is

set up between the service end-points and the access control bandwidth limits are modified. When the service is over, the tunnel is automatically removed and access control limits are reinstated for the customer. All these operations are performed automatically without human intervention: the customer is 'self-provisioning' the request on the network.

At all times, the BoD system controls the available bandwidth and rejects user demands that exceed the available capacity for the requested time period. In this way QoS is guaranteed for all services at all times.

The system uses different hybrid solvers:

- the Route Generator module finds routes in the network which satisfy delay and capacity limits
- the Commitment Solver compares the committed bandwidth in the network with current and historical traffic measurements and determines the available capacity for additional services.

These solvers are implemented in ECLiPSe, our constraint-programming

environment, and use a combination of constraint, local search and mathematical programming techniques. The hybridisation allows us to find near-optimal solutions in a limited time, and to scale the algorithms even for large networks.

The system was developed jointly by Parc Technologies and IC-Parc, Imperial College London for an industrial partner, and has been implemented on a global MPLS network. Parc Technologies is a start-up company developing decision-support systems for network problems. IC Parc is a research department in the faculty of Engineering, Imperial College London. The Parc Bandwidth on Demand Solution is only one of several constraint-based systems that have been developed by Parc Technologies in the networking area in recent years.

Link:

<http://www.parc-technologies.com>

Please contact:

Helmut Simonis, Parc Technologies Ltd, UK
Tel: +44 207 594 8463
E-mail:
helmut.simonis@parc-technologies.com

DesParO - A Design Parameter Optimisation Toolbox using an Iterative Kriging Algorithm

by Christof Bäuerle, Clemens-August Thole and Ulrich Trottenberg

DesParO is an optimisation toolbox designed for industrial simulation. The toolbox contains a collection of efficient algorithms for computer-based optimisation and can easily be adapted to any industrial simulation code. DesParO was especially designed for computationally expensive simulation codes and supports parallel computation.

Nowadays simulation programs are widely used in the car industry to optimise production processes or to test the behaviour of a car before a real prototype is built or real tests are performed. Simulation helps to reduce the number of real prototypes and the time between model changes compared to real tests. Crash simulation is one of the most time-consuming simulation tasks and is therefore computationally expensive. Many simulation runs with different parameter values are required to optimise construction parameters, and the results have to be analysed by an engineer. In this case the use of a numerical optimisation toolbox, which analyses different parameter constellations, can support the development process significantly. The optimisation toolbox uses the PAM CRASH simulation program to calculate the crash properties of a car model. For each given set of parameters the simulation results have to be evaluated by an objective function. This article focuses on a box beam test example which demonstrates the applicability of the toolbox to crash simulation problems.

The definition of an objective function which evaluates the crash simulations can be difficult for technical applications. In the case of the box beam example the objective function is constructed as a comparison between a reference crash and a crash of a model with modified material parameters. The aim of this procedure is to define an optimisation problem with a well-known solution to test the toolbox and the optimisation algorithms. In the reference case the acceleration of a box beam with given Youngs modulus and Poisson ratio is calculated and the objective function is defined as the integral of the quadratic differences to the acceleration of a model

with modified parameters over the time. This integral can be approximated by a sum over discrete time steps:

$$f = \int (a_r(t) - a(t))^2 dt \approx \sum (a_r(t_i) - a(t_i))^2 \Delta t_i$$

Thus the minimum of this objective function is reached if the material parameters of the modified model reach the values of the reference case.

Figure 2 shows a plot of the objective function versus change of the Youngs modulus. The minimum of this function is at 210 kN/mm² and is identical with the reference case. Close to its minimum

the function is smooth but on the left side the behaviour is rather chaotic. Under realistic conditions, almost all properties of crash experiments are subject to scatter, although it is desirable to avoid the scatter of a model as much as possible. Numerical properties of the simulation codes as well as certain features of the crash model may be responsible for these in-stabilities. Typical sources of scatter which are caused by the model are buckling and contact under an angle of 90°. Further investigation of this effect with the box beam example indicated that the scatter results from an instability between two




Figure 1: Crash example box beam.

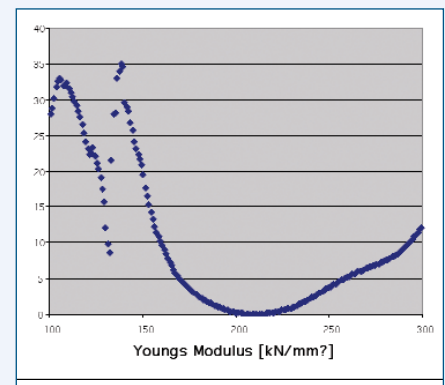


Figure 2: Scan of the objective function for the test problem.




Figure 3: One possible mode of the crash simulation.

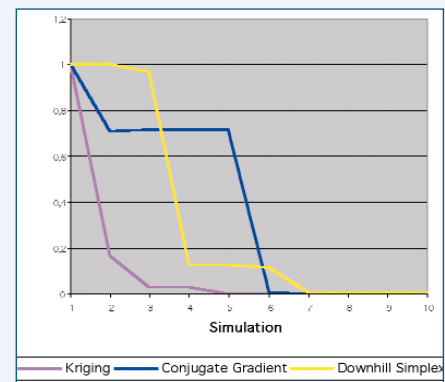


Figure 4: Convergence of the different optimisation algorithms.

different modes. Figure 3 shows one such mode. In this case the upper edge on the right side is shifted outward and the lower edge is shifted inward. In the second mode of the model both edges were displaced outside the box beam.

Scatter of a model is a difficult condition for optimisation and an engineer tries to avoid instabilities of a construction. In the test case of the box beam example it turned out to be possible to stabilize the model by a small modification of the construction. In particular, if both edges of the original model are bent to one side it is possible to suppress the asymmetric mode during the crash simulation and the scatter of the objective function can be reduced significantly. This procedure now leads to a well defined optimisation problem.

For the solution of nonlinear problems a wide literature on optimisation strategies is available. Popular strategies are conju-

gate gradient, downhill simplex, and surrogate models. An important characteristic of the optimisation of crash simulations is that it is usually very time-consuming for a realistic car model. Optimisation strategies based on surrogate models are especially suited for optimisation under such conditions. In the first step design of experiment methods are used to choose parameter values, for which a simulation has to be executed. The values of the objective function for these parameters can now be used to construct a surrogate function as an interpolation of the real problem. The evaluation of the surrogate function is very fast and it is much easier to search for the minimum of this interpolation. During the next steps the design of experiment points can be refined iteratively nearby the minimum of the surrogate function. In Kriging the surrogate model consists of a sum of basis functions, very often of Gauss type. Figure 4 shows the relative convergence of different optimisation

algorithms at the stabilized box beam example which was described in this article. Although the Kriging algorithm converges very fast during the first iterations at this problem, the optimisation strategy does not work so reliably if the objective function is not smooth enough.

Future work will concentrate on the aspect of improving the robustness of the surrogate models by replacing the design of experiment methods by pattern search strategies. Furthermore it is planned to extend the surrogate models to unstable crash constructions. In this case surrogate functions are used for approximation (instead of interpolation) of the objective function.

Links:

<http://www.scai.fraunhofer.de/35.0.html>
<http://www.auto-opt.de>

Please contact:

Christof Bäuerle, Fraunhofer SCAI
 Tel: +49 2241 14 27 94
 E-mail: baeuerle@scai.fhg.de

On-line Testing of the Reactor Protection System in the Paks Nuclear Power Plant

by István Varga, Tamás Bartha and Alexandros Soumelidis

The most important issue for safety-critical supervisory systems is the correct operation of every component. The Systems and Control Laboratory (SCL) of SZTAKI has developed a new validation and testing concept to ensure the correct functionality of safety elements. Three implementations have been completed and are regularly used in the Paks Nuclear Power Plant (Hungary).

Instrumentation and control (I&C) systems are used in safety-critical applications to monitor issues such as operational safety and service availability. They are gradually being replaced by software-based digital systems, and in order to enhance reliability in the design and operation phases of these software-based I&C systems, a number of techniques must be applied. One of the best-known of these is periodic testing during normal operation.

We have created a testing methodology for safety-critical supervisory systems, which has been implemented at the Paks Nuclear Power Plant located in Hungary. Each of four reactor units has a dedicated

Reactor Protection System (RPS) for the intervention and safe shut-down of the unit in an emergency situation. This system contains three redundancies, also called 'trains'. An RPS, like all physical systems, is also susceptible to faults, including sensor failures and component hardware faults. Naturally, each RPS has a fault-tolerant, highly redundant hardware and software architecture. However, faults can cause latent errors that may become activated only in emergency conditions. To eliminate these and ensure reliable operation in all situations, an RPS must be periodically tested during normal operation by special test equipment.

The Test Concept

It is necessary that the test equipment be able to create experimental scenarios of input signals in one part of the protection system, without affecting the operation of the other parts or the RPS as a whole. While testing a certain train in a 2/3 redundancy system, only two trains remain in normal operation. This situation is very sensitive, since an error occurring in either of the two operating redundancies could initiate an emergency protection action. The train to be tested must be selected by a corresponding 'Test Enable' (TE) signal, while the TE signals of other trains are simultaneously disabled. During the test,

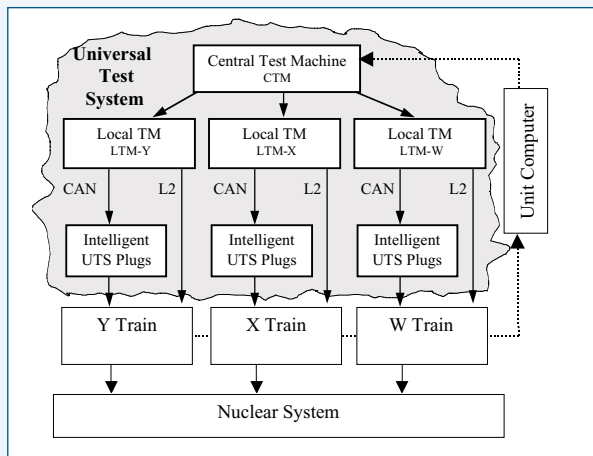


Figure 1: Structure of the Universal Test System.



Figure 2: UTS plugs.

the process parameters are set in such a way that safety actions are initiated.

The test execution makes the 2/3 redundancy system sensitive to errors. For this reason, the two trains not being tested are constantly monitored for errors during the test. If an error is detected in these systems, the test stop criterion disables the TE signal of the train being tested. This causes the test to cease immediately, and the train returns to normal operation. This ensures that two fault-free trains are always available for correct voting, and thus we avoid the activation of unnecessary EP actions.

The Test Equipment

The principles shown above have been implemented in the RPS of the Paks Nuclear Power Plant. We created three realisations of the test concept: test machines for the start-up and operational phases of the reactor (the Start-up Test

Machine and Periodic Test Machine respectively), and an improved test system, which includes the complete functionality of the earlier machines (Universal Test System).

The periodic testing concept of the RPS system suggests a stand-alone Test Machine, which is temporarily connected to test inputs from the system in order to perform test procedures. However, a Periodic Test Machine designed and built on this basis proved inconvenient and difficult to use in practice. Therefore, the revised test concept includes a distributed, general-purpose test system called the Universal Test System, or UTS for short.

Universal Test System

The basis of the distributed test equipment is an RS485-type industrial bus system (CAN), which provides the interface to the task-specific ‘active’ test

plugs (analogue test plug, binary test plug, control plug etc; see Figure 2). The microprocessor-controlled analogue and binary active plugs contain all the devices necessary to perform a full test of the analogue and binary input modules. The active control plug contains transceiver and receiver components to send and receive control signals to/from the RPS system. The design of the active test plugs, the microcomputers and the RS485 bus interface requires implementation in high density SMD technology.

The active test plugs are powered by the RPS system. When all the safety requirements are equally satisfied, galvanic isolation must only be realised on the CAN interface. The analogue and binary plugs can be applied simultaneously to multiple input modules, even to a whole cabinet or a complete train. An industrial computer (Local Test Machine) controls the active test plugs using a CAN controller card. The Central Test Machine (CTM) supervises the three Local Test Machines, each of which is dedicated to a single RPS train, as shown in Figure 1.

The Universal Test System is in regular use at the Paks NPP, and has helped to improve the reliability of the RPS system, while reducing the time, and therefore cost, dedicated to testing.

Link:
<http://www.sztaki.hu/scl>

Please contact:
 Tamás Bartha, István Varga, SZTAKI
 Tel: +36 1 279 6227
 E-mail: bartha@sztaki.hu, ivarga@sztaki.hu

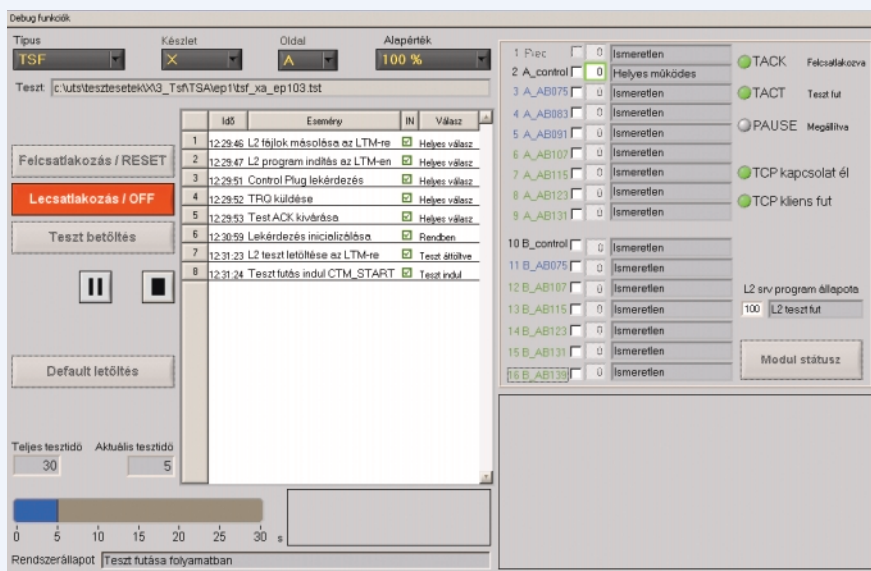


Figure 3: The CTM software during a test.

Process Systems: Theory and Applications from Different Aspects

by Katalin Hangos and Gábor Szederkényi

Combining a priori physical and engineering knowledge with theories of artificial intelligence, discrete event systems and modern nonlinear control has led to new theoretical results and successful applications in the field of modelling, analysis and control of process systems at the Process Control Research Group at SZTAKI.

A rough definition for process systems is that the operation of such systems is based on the laws of thermodynamics. Well-known examples of process systems include heat exchangers, chemical reactors, distillation columns and wastewater treatment systems. The Process Control Research Group is part of the Systems and Control Laboratory of SZTAKI, and studies the modelling, analysis and control of process systems.

Process Modelling and Model Analysis

Process models for control and diagnostic purposes most often take the form of differential-algebraic equations (DAEs), with the conservation balances being the differential part and the constitutive equations the algebraic part. A number of hidden relationships exist between the equations and their terms, as dictated by the laws of thermodynamics, unit operation, reaction kinetics, heat and mass transfer. Together with these relationships, the equations can be seen as structured texts (like computer

programs) with a syntax and semantics defined implicitly by the underlying laws of thermodynamics and chemical engineering. This approach enables us to use formal methods of computer science to construct, verify, analyse and simplify process models in a rigorous and automated way. Our latest results in this field are an assumption-driven model-building procedure, a systematic model-simplification method and an assumption-retrieval procedure for the determination of simplification assumptions of two related process models. Furthermore, the prototype of an assumption-driven model editor has been developed in Prolog language. Cooperating partners in this research area are the Department of Computer Science at the University of Veszprém in Hungary and the Computer Aided Process Engineering Centre in the Department of Chemical Engineering at the University of Queensland in Australia. Several years of joint work resulted in a book entitled 'Process Modelling and Model Analysis', written

by Katalin Hangos and Ian Cameron and published by Academic Press in 2001.

Intelligent and Discrete Process Control

Intelligent and discrete process control and diagnostic methods apply so-called discrete event system models, which are currently one of the hottest research topics in systems and control theory. Such systems are widely found in practical applications and are characterised by the presence of discrete valued variables: state, input and output variables with a finite number of possible values. At present, our research in this field is focussed on the hierarchical decomposition and analysis of process models using hierarchical coloured Petri Nets. Our previous results on the invariance inheritance in Petri Nets and the application of Artificial Intelligence methods in generating operating procedures serve as a starting point for this work. Recently, we successfully utilised hierarchical coloured Petri Nets for multiscale-modelling, deadlock-detection and fault




Figure 1: Reachability surface in the state space of a simple fed-batch fermentation process.

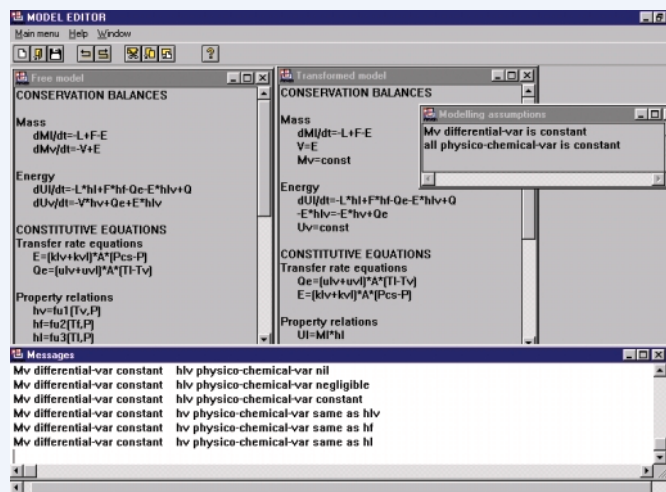


Figure 2: Screenshot from the assumption-driven model editor software (analysis of model transformations).

diagnosis in an industrial granulator system. The simulation results run on the validated system model are promising, making the next short-term goal to implement the developed techniques on the real system.

In 2001, Kluwer Academic Publishers published a book in this field with the title 'Intelligent Control Systems (An Introduction with Examples)', written by Katalin Hangos, Rozália Lakner and Miklós Gerzson.

Nonlinear Process Control

Nonlinear system theory and the nonlinear control and diagnosis methods based thereon form one of the most challenging and fast-developing areas in modern systems and control theory. Its significance is explained by the fact that most engineering systems (and process systems in particular) show nonlinear behaviour, with a few exceptions of continuously operated systems near a well-defined and stable operation point.

Apart from its practical importance, nonlinear systems and control theory requires the understanding and application of advanced mathematical concepts and tools from the areas of nonlinear differential equations, functional analysis, differential manifolds and algebra and symbolic computations. It turns out, however, that the principles of the engineering discipline governing the behaviour of the nonlinear system play a key role in the understanding, analysis and control of engineering systems. Thus the use of engineering knowledge determines certain structural properties of a nonlinear system, and this can make it feasible or even easy to handle otherwise very difficult problems. This type of interdisciplinary approach to control problems is commonly called a 'grey-box' approach. Recent research results show that this approach can be successfully applied for the nonlinear analysis of control-related features such as local and global stability, reachability, observability, passivity and zero dynamics. These analysis results together with physical and engineering knowledge allow us to understand the dynamic behaviour of process systems in more depth, and to select and design nonlinear

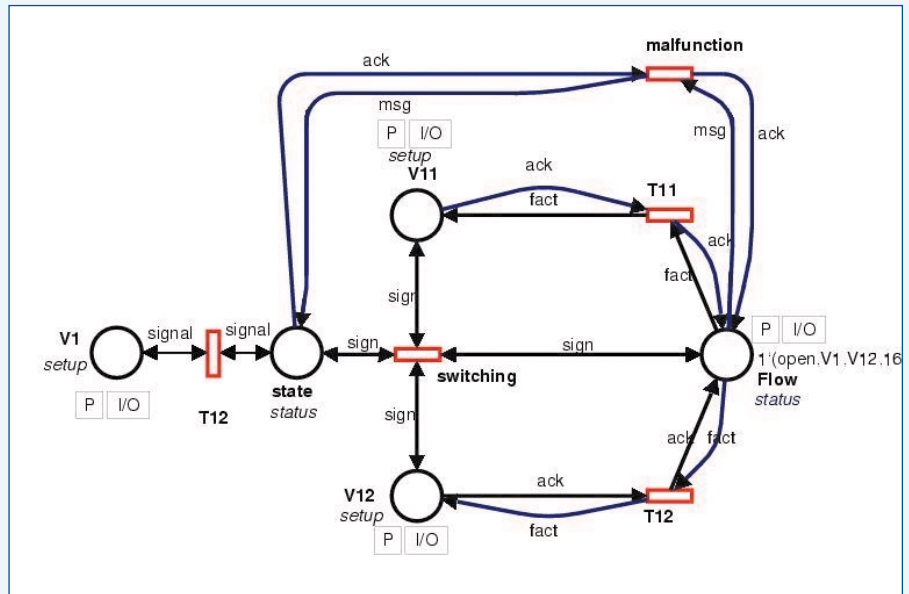


Figure 3: Coloured Petri Net model of a valve sub-system.

controller structures with the properties we require.

In the field of nonlinear mechanical systems, the Hamiltonian description of these systems gives rise to their special characterisation and the development of their very powerful passivity-based control. The analogy with the Hamiltonian description of mechanical systems has led to the Hamiltonian description of process systems.

The developed methods were successfully tested on a variety of process systems:

- fermentation processes with highly nonlinear reaction kinetics operated in continuous or fed-batch mode
- a low-power nonlinear gas-turbine (in cooperation with the Department of Aircraft and Ships at the Budapest University of Technology and Economics and the János Bolyai Military Technical College Faculty of the Miklós Zrínyi National Defence University)
- a nonlinear electro-pneumatic valve within the brake system of trucks (in cooperation with the Research and Development Center of Knorr-Bremse Brake-systems Ltd. in Budapest).

Our most important partner in this area of activity is the Systems and Control Laboratory of SZTAKI, and a significant

result of this cooperation is another book, entitled 'Analysis and Control of Nonlinear Process Systems' and written by Katalin Hangos, József Bokor and Gábor Szederkényi, which is to be published by Springer-Verlag in early 2004.

Links:

<http://www.sztaki.hu/scl>
<http://daedalus.scl.sztaki.hu>

Please contact:

Katalin Hangos, SZTAKI
 Tel: +36 1 279 6101
 E-mail: hangos@scl.sztaki.hu

Multi-Agent System Technology in a Port Container Terminal Automation

by Vicent J. Botti

Scientists at Universidad Politécnic de Valencia have been working on a solution to a port container terminal management, specifically the automatic allocation of containers. Given its inherent complexity, we propose that multi-agent systems are the most suitable method for addressing this problem.

Agent/multi-agent systems have become an important field within artificial intelligence research. They have found a number of applications, including control processes, mobile robots, air-traffic management and intelligent information retrieval.

Here we present a multi-agent system for solving the automatic allocation problem in a container terminal. The operations carried out in such terminals are among the most complex tasks in the transport industry. This is due to three factors: the diverse range of entities performing operations, their interaction with a dynamic environment, and the distributed nature of the problem, in which independently functioning systems directly affect each other through their actions.

Existing centralised and sequential applications for container terminal management are now insufficiently flexible to respond to changing management styles and the highly dynamic variations in loading/unloading requirements. Traditionally, the entire terminal is controlled by centralised software, which limits the expansion and reconfiguration capabilities of the system. Using hierarchical organisation forces the grouping of resources into permanent,

tightly coupled sub-groups, within which information is processed sequentially by a centralised software supervisor. This may result in much of the system being shut down by a single point of failure, as well as contributing to plan fragility and increased response overheads.

The multi-agent system model seems an adequate framework to overcome such problems through the design and development of an application that is flexible, adaptable, versatile and robust enough for the efficient management of a container terminal.

It is critical that once in port, the turnaround of a cargo ship be as fast as possible. An average cargo liner spends 60% of its time in port, and a cost in the order of US\$1000 per hour is associated with this. Consequently, the container allocation process must be directed towards minimising cargo ship stowage time. This is the main objective of the optimisation of the global performance allocation process.

Container terminal management is a very complex system, and traditional solutions decompose the problem into several sub-problems, each representing one particular aspect. While the set of

operations to be conducted in the terminal is very extensive, existing approaches share some common features:

- the Marine Side Interface focuses on the loading and unloading of containers. Normally two or three gantry cranes (GC) are used to move containers for each ship
- the Transfer System transfers containers between the apron and the container storage yard. Yard trucks (YT) perform transports within the terminal. Transtainers are used to pick up or to put down a container on the storage area of the yard (see Figure 1)
- the Container Storage System allocates and controls containers in the yard (see Figure 2)
- the Land Side Interface handles interactions with the land transportation modes.

System Architecture

Figure 3 shows the system architecture, in which we divide the problem into sub-problems, each of which is solved by a specific agent. These agents are mainly characterised by their independence from the rest of the system elements. They are able to coordinate and to communicate decisions to the rest of the system. Communication between agents is done by means of asynchronous




Figure 1: Yard map.




Figure 2: Transtainer.

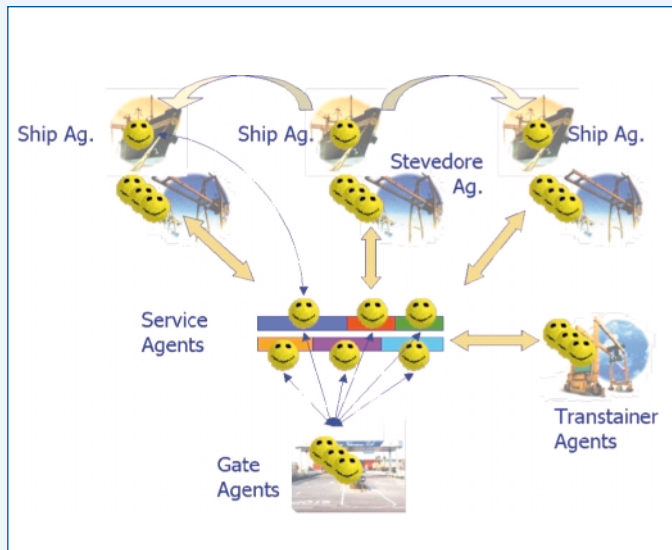


Figure 3: System Architecture.

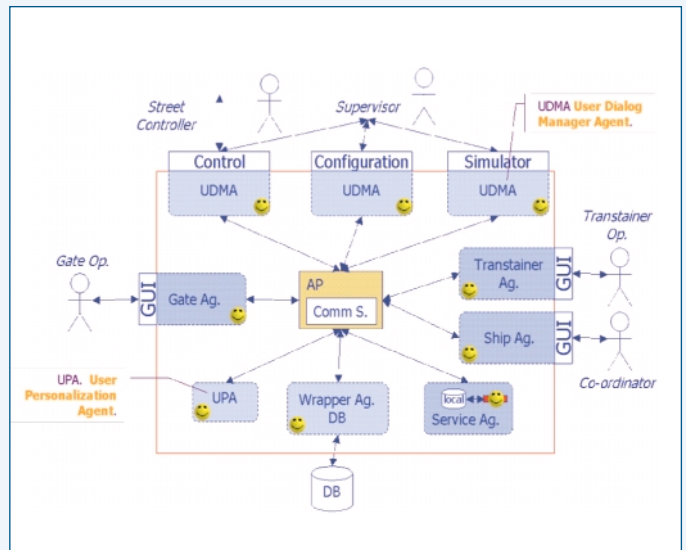


Figure 4: System Interactions.

messages, which are based upon the FIPA-ACL standard. The proposed distributed approach enhances flexibility, efficiency and robustness. Five agent classes can be found in this system:

- *ship agents* determine the ships' loading/unloading sequence
- *stevedore agents* manage the ships' loading/unloading process
- *service agents* distribute the containers in the port terminal
- *transtainer agents* optimise the use of these machines
- *gate agents* interact with the land transport (I/O of containers by land).

Ship Agents

In response to the arrival of a ship (ship agent creation event), the system creates a new ship agent for this ship and its load profile. Its goals are to minimise gantry crane idle time, the ship's loading/unloading time, and the derived costs from the stowage process. This work is closely related to that of the stevedore agents, with which the ship agent must coordinate.

Each ship agent faces a scheduling problem in which a set of resources (the cranes) must be assigned to the different operations (loading/unloading of containers), thereby establishing a resource use time (loading/unloading time). This requires that the various ship agents active at any time must coordinate with each other to minimise clashes between assigned cranes.

Stevedore Agents

For any given gantry crane, stevedore agents try to obtain the most appropriate scheduling in order to manage container stowage. The agent must know the gantry crane loading/unloading sequence, which yard trucks are assigned to this crane and the positions of the various containers within the terminal (this information is provided by the relevant service agents). The agent therefore coordinates with active ship agents and service agents, and attempts to minimise both empty movements of the machinery employed and the number of machines necessary for the internal transfer.

Service Agents

The terminal is divided into services, with each being assigned specific stacking ranges. The main goal of a service agent is to determine the appropriate allocation for containers arriving in the terminal from a specific service, and the most suitable configuration for that portion of the yard controlled by the agent. To do this, the agent must know the yard map of its assigned portion, the container characteristics (type, length, weight, destination port, ship) and the stacking factor. The agent must also coordinate with other service agents in order to resolve conflicts, such as reallocation of containers where their assigned stacks are full.

In solving the configuration problem, service agents must maximise stacking density in their yard portions. The

service agent launches this process automatically when it considers it to be necessary, based on criteria such as time, stack allocation conflicts (slots without use or cargo ships that run out of reserved areas), low stacking density etc.

Transtainer Agents

Each transtainer is modelled as an autonomous agent whose goal is to perform container-stacking operations efficiently. A transtainer agent must minimise the transtainer's empty movements. To do this, it obtains the most efficient sequence for moving the container to or from its correct position in the yard. Each agent waits for stacking requests from the various service agents, which inform the transtainer agent of the location of containers to be loaded to vessels or external trucks, or where containers being unloaded from vessels or trucks are to be placed.

Gate Agents

A gate agent controls the arrival or departure of a container by land. The agent manages the assigned terminal gate, informing the corresponding service agent of the arrival of new containers (in order to store them) and of the arrival of trucks (in order to retire containers from the yard). For instance, when a container arrives, the gate agent checks the accuracy of its data, and if the data is correct, asks for a container location from the service agent responsible for the service to which the container is assigned. Once the location is known,

this is communicated to the truck, which delivers it to the appropriate stack. A similar process occurs for containers leaving the terminal.

Implementation

Internal and external interactions (between the various agents and with system users respectively) are shown in Figure 4. FIPA standards have been followed as far as possible in imple-

menting this prototype. In doing so, a set of auxiliary agents has been created:

- wrapper agents provide access to the database and communicate with external software
- the UDMA (User Dialogue Management Agent) is an interface agent with human users
- the UPA (User Personalisation Agent) manages the explicit profiles and preferences of registered human users.

There is also an agent platform (AP) to support the entire architecture. It provides the following services: yellow pages, white pages, a communication channel and an agent platform security manager.

Please contact:

Vicent J. Botti, Universidad Politécnica de Valencia/SpaRCIM
E-mail: vbotti@dsic.upv.es

Contribution to Fault Detection and Diagnosis: Mixture-based Modelling

by Tatiana V. Guy, Dani Juriš, Miroslav Kárny and Andrej Rakar

A multi-model approach complemented by probabilistic mixture modelling has been applied to solve fault detection and isolation problem. The resulting algorithm provides efficient tool for process condition monitoring.

Complexity of modern manufacturing and process industries emphasises the importance of the fault detection and isolation for the purpose of process condition monitoring. Unlike occasional inspections, periodic checks, etc, on-line condition monitoring seems to be the most promising way to decrease the number of unanticipated system shut-downs while optimising the production costs. Extension of recently developed probabilistic mixture-based approach to the domain of systems' fault detection and isolation provides an adequate support to solution of process monitoring problem. Performance of the approach was verified on a real process of gas conditioning.

Within an EC IST-99-12058 project ProDaCTool, a powerful generic tool for describing and optimising complex uncertain multi-dimensional systems by multiple models was developed and successfully applied. Hypothesis on existence of several operational regimes - normal as well as faulty - makes multi-model approach a natural framework for solving the fault detection and isolation (FDI) problems. This has motivated us to test its applicability to FDI domain.

The whole framework relies on a probabilistic mixture modelling that provides suitable description of non-linear dynamic systems. Each of the operation modes as well as each of the faulty states is represented via probabilistic model. A convex probability density functions $\hat{\cdot}$ called components - describing respective normal dynamic regressions, each associated within some operation mode, models the process data. Quasi-Bayes estimation provides the model of the supervised process. Provided rich process data exist, the model learned represents all the operation states occurred in the past, each associated with the corresponding component. The probability of occurrence of a particular operational state is expressed by the weight of the respective component. During on-line operation, the learned components' weights are updated using available process measurements and exploited for effective estimation of the actual operation state. A proper recommendation is then generated in a timely fashion. Effectiveness of the outlined approach has been verified on a gas-conditioning unit, a part of an industrial-scale pilot installation in Josef Stefan Institute for treatment of technological

wastewater. Real-time diagnostic results show that each component correctly represents faulty states of the system. Apart from short transient periods, during changes of process states, all faults were successfully detected and uniquely isolated. It is important to note that the presented diagnostic approach relies on availability of rich past data. These should also include all faulty states that have to be diagnosed. Unfortunately this is not always the case in real applications, especially if a diagnosis system is designed from scratch. The on-line applicability of the quasi-Bayes estimation allows to design adaptive FDI version. The available self-learning capability will enable to consider newly met faulty states.

Please contact:

Tatiana V. Guy, Institute Information Theory and Automation, Academy of Sciences of the Czech Republic /CRCIM
Tel: +420 2 6605 2254
E-mail: guy@utia.cas.cz

Data Mining Applied in Anaerobic Wastewater Treatment

by Simon Lambert

The research and development project TELEMAC is creating an adaptable, customisable system for the monitoring and control of anaerobic wastewater treatment plants. One of the techniques being employed is data mining — the extraction of useful knowledge from data using a variety of techniques.

Coordinated by ERCIM, TELEMAC is a project within the European IST program that is developing innovative approaches to managing anaerobic wastewater treatment plants, with particular application to the wine industry (see also ERCIM News No. 48). One of the features of the project is the integration of a variety of approaches, including soft sensors, fault detection and isolation, and remote monitoring and access of multiple plants. The project has now been running for two years and has made significant progress.



Anaerobic digester at industrial scale (Sauza, Mexico).

Within TELEMAC, CCLRC has been working specifically on data mining. This is an important approach, since data, such as pH, temperature, and more advanced measurements such as volatile fatty acids (VFA), is constantly being accumulated. Data mining opens up the prospect of learning from this data in order to manage plants better. A number of possibilities are being studied:

- developing models or rules that help to predict dangerous conditions on the plant from trends in sensor readings
- detecting faulty sensors through inconsistent sets of readings
- partially substituting for expensive sensors by combining readings from more commonly available sensors.

Within the TELEMAC project, data is available from a number of plants of differing types and sizes. These range from large industrial waste-processing plants, through pilot-scale plants with a full range of instrumentation, to small laboratory-scale set-ups used for running specific experiments.

Data mining is often regarded as one part of the broader problem of knowledge discovery. Knowledge Discovery in Databases (KDD) is defined as ‘a non-

trivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data’ and data mining as ‘exploration and analysis, by automatic or semi-automatic means, of large quantities of data in order to discover meaningful patterns’ (Kumar, V. & Joshi, M. [1999], Tutorial on high-performance data mining, at <http://www-users.cs.umn.edu/~mjoshi/hpdmntut/>, University of Minnesota). Although much of the activity is data mining, the goal is usually knowledge discovery. Data mining is not simply a matter of running algorithms for rule induction or neural networks; a considerable amount of preliminary work is usually required, and TELEMAC is no exception. The overall process includes stages such as data selection, data cleaning (for example, dealing with missing values or outliers), data reduction or enrichment, data preparation, the data mining itself, and reporting (through visualisation, statistics etc).

The software being used at CCLRC for data mining is Clementine, a product available from SPSS. Clementine was developed in Europe as a commercial, general-purpose data mining tool and was adopted by the Business and Information Technology Department of

CCLRC in 1998 for scientific applications such as the EC-funded project DECAIR (<http://www-air.inria.fr/decair/>). Like a number of visualisation systems (for example AVS, Nag Iris Explorer, IBM Open DX), it presents a visual interface that allows a user to connect modules together, allowing data to flow from one end to the other. Visualisation software such as XMDV is also being used.

Initial work focussed on the use of data mining to recover a known model from synthetic data, thereby giving prima facie assurance of the applicability of the approach. Fitting was generated using a simultaneous prediction of three key variables. A Pruned Neural Net modelling tool was employed. The data was split into a test set and a training/validation set. The training/validation set is randomly split 50/50. The reason for this three-way split is to avoid over-training the neural net — a common difficulty faced in practice. It has proved possible to reproduce the synthetic data to a good level of accuracy.

Exploratory work has also been done on confidence and prediction intervals for TELEMAC data. This refers to the distinction between the accuracy of the fit itself (the regression) and in the predictions arising from its use. These are different, since an individual value will have extra variability due to noise.

Simple modelling of target variables for some of the industrial data is also under way. The objective is to determine whether a neural net model constructed in one time period could be used to make predictions of a target variable in another time period. The technique used is the Extended Pruning approach with a measurement of VFA as the target, and it

gave 84% accuracy when applied to data from another time period from a particular plant.

Techniques of rule induction are being applied to estimate values of sensor readings based on more easily obtained values, and to determine how reliable the models so developed remain over time. Rules are generated in forms such as the following:

Variable X falls in a particular range of high values if variable Y falls within a particular range of low values – this rule scores (N, p).

Here N and p are indications of the degree of satisfaction of the rule. Rule induction has the advantage that the rules generated are often meaningful to human

domain experts, and can be critiqued and validated by them.

A key question for future work is the handling of the evolution of the state of the plant over time. This will begin with the approach of Dorffner to neural networks for time series processing, in which a model is built up from training records that contain both prior time values and current values. This technique is often known as lagging.

A further important question is the place of data mining in the final TELEMAT system. The project as a whole is now considering the form of the final TELEMAT system and its tools, but at the time of writing it seems that data mining will operate as a component at the

Telecontrol Centre (responsible for monitoring several plants), and will run periodically to update the knowledge base. It remains an open question as to how to detect when it is necessary to re-run the data mining algorithms for a particular plant.

Link:

<http://www.ercim.org/telemac/>

Please contact:

Simon Lambert, CCLRC
Responsible for data mining work
Tel: +44 1235 445716
E-mail: S.C.Lambert@rl.ac.uk

Bruno Le Dantec, ERCIM
Administrative coordinator
Tel: +33 4 9238 5013
E-mail: bruno.le_dantec@ercim.org

Olivier Bernard, INRIA
Scientific coordinator
Tel: +33 4 9238 7785
E-mail: Olivier.Bernard@sophia.inria.fr

Integrated Multivariate Statistical Process Control and Condition Monitoring for Fed-batch Fermentation Processes

by Hongwei Zhang

Multivariate Statistical Process Control is seen by many as the key to real-time process monitoring which results in significant improvements in operational safety and efficiency. Members of the Control Engineering Research Group (CERG) at Manchester School of Engineering are developing MSPC technologies, including Principal Component Analysis (PCA), Partial Least Squares (PLS) and Kernel Density Estimation (KDE), and combining them to form an integrated robust process control and monitoring package. Recently, research has been carried out at CERG with particular interests in the applications to fed-batch fermentation processes.

As fermentation processes are responsible for the production of high value added substances such as enzymes, antibiotics and vitamins, they have great importance in the biotechnology industry. From the control engineer's point of view, the fed-batch processes present the greatest challenge. At the same time, there is an enormous economic incentive behind the work on control, optimisation and condition monitoring of fed-batch fermentation processes.

To achieve high performance operation, the optimisation of the primary factors influencing fed-batch fermentation processes is an important task. The

ability of optimal control techniques to assist in this optimisation has been investigated. Optimal control techniques rely upon an accurate model of the process and for many years mechanistic models have been used to develop optimal control strategies for fed-batch processes. However, mechanistic models of fed-batch processes are usually very difficult to develop due to the complexity and nonlinear nature of the processes. An attractive alternative to developing a mechanistic model of a process is to build an empirical model using input-output data collected from the process. Among all the existing methods, the multivariate statistical process control techniques, such as PCA

and PLS have been proved to be very promising approaches for application to fed-batch fermentation systems. Multiway Partial Least Squares (MPLS) is used to identify empirical models for fed-batch fermentation processes. The objective function of the optimal control problem for a batch or fed-batch process needs to reflect the performance of the fermenter. Maximisation of yield, selectivity or conversion and minimisation of batch cycle time are some examples of possible objectives. As the optimal control of fed-batch fermentation processes is a challenging dynamic optimisation problem, it is usually difficult to solve because of the nonlinear system dynamics and the constraints on the

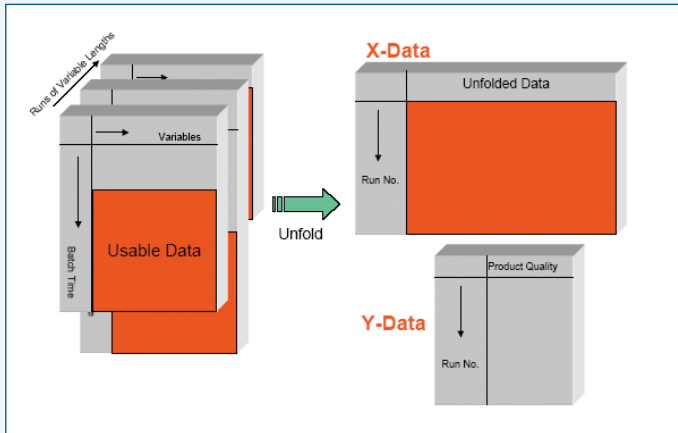


Figure 1:
Data Unfolding.

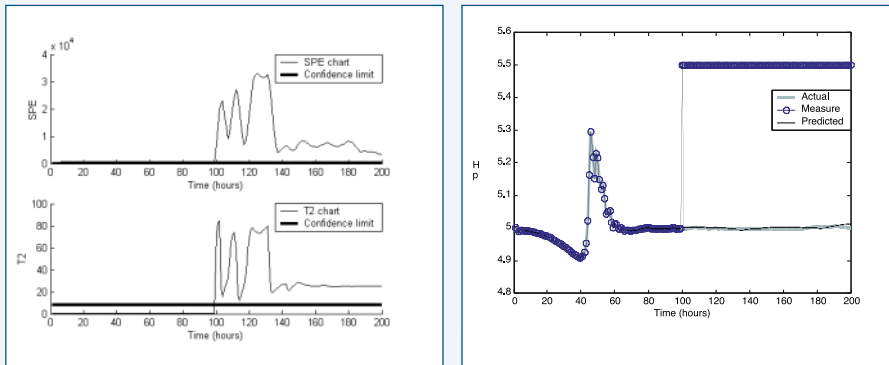


Figure 2: Detection of a pH sensor fault and pH measurement inferred by the PLS model.

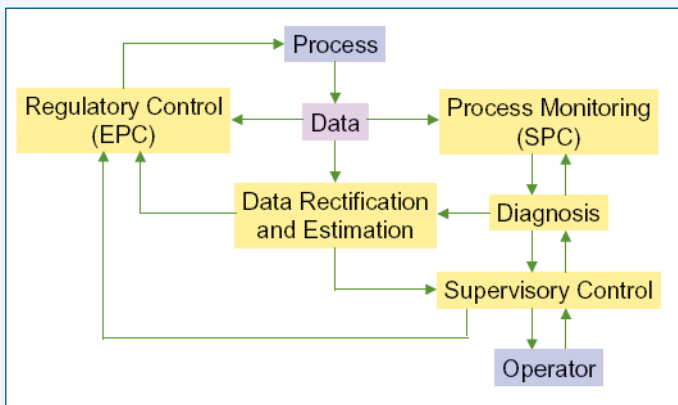


Figure 3:
Integrated Process Control and Condition Monitoring.

control and state variables. Typical choices for the objective function in the optimisation of a fed-batch fermentation process include the maximisation of either biomass or metabolite production with respect to the substrate feed rate. Based on the techniques of MPLS and FGA, members at CERG have developed optimal control algorithms that are suitable for application to batch processes. The algorithm utilise linear models of the process that are identified using the multi-way partial least squares technique. The linear models are used within optimal control law to regulate the productivity of the batch by manipulating the substrates that are fed into the

fermentation vessel. The proposed control algorithms are successfully applied to different fed-batch fermentation process and their performances compared with alternative control algorithms that have been applied to the applications in the past. The proposed controllers are found to compare favourably with these alternative algorithms.

As the objective of control systems for batch and fed-batch processes is always to ensure high quality of the products. A further issue in the successful operation of fed-batch fermentation processes is the early and accurate detection of fault

conditions, such as sensor failures or drifts that may occur. The early detection of fault conditions is of great benefit in fed-batch fermentation processes since the earlier that a fault can be detected and acted upon, the lower its impact will be on the process. In some situations this can be critical, for example, a drift on a pH sensor could have catastrophic results on biomass growth if this measurement is used within a feedback control scheme. PCA and PLS have also been seen as very promising approaches for fault detection and isolation. In previous studies condition monitoring and control of fed-batch fermentation processes have been viewed as independent problems. In the research carried out at CERG, these problems are considered together and integrated fault detection and process control schemes are being developed. These schemes rely heavily on the successful development of a PLS model to provide soft-sensing capabilities in a fed-batch fermentation processes. In research implemented at CERG, it is demonstrated through several different case studies that PLS models can be used, not only in a predictive control framework to regulate the growth of biomass within the fermenter, but to also provide fault detection and isolation capabilities. It is further shown from simulations that the integration of the predictive controller and fault detection capabilities provides a useful diagnostic tool for the control system.

This research so far has seen some very good results in various simulations, efforts now are being made to verify these techniques through lab experiments. Future research also includes the integration of artificial intelligence methods such as expert systems and the multivariate statistical process control methods. This project is funded by EPSRC (Grant Number GR/N24858).

Links:

<http://travis.eng.man.ac.uk/>
<http://www.newi.ac.uk/zhangh/>

Please contact:

Barry Lennox, University of Manchester, UK
Tel: +44 161 275 4324
E-mail: Barry.Lennox@man.ac.uk

Hongwei Zhang, Northeast Wales Institute of Higher Education, UK
Tel: +44 1978 293 392
E-mail: H.Zhang@newi.ac.uk

Saving Energy

by Claude Lemaréchal

After fifteen years of collaboration, EdF, the French Electricity Board and INRIA are finalising an efficient software which will optimise the daily production of electricity in France.

To meet a demand that oscillates between 40 and 60 gigawatts, EdF operates about 150 thermal units (nuclear and conventional) and 50 hydro-valleys. Since electricity can hardly be stored, production must closely follow demand. Changing the operation mode of a unit however, especially a thermal one, is difficult and costly. As a result, optimising the schedules is a critical problem, involving large amounts of money. It is also a tough nut to crack, and is often referred to as the unit-commitment problem: large-scale, mixed-integer, heterogeneous. Its complexity can be illustrated by a superficial description of the production units.

A thermal unit must produce a certain minimum power (positive). This introduces a combinatorial aspect into its operation - the unit is 'on' or 'off'. This is not the only one, however: when switched off or on, the unit must respectively stay off or maintain its production level, for a minimum period of time. Also, when switching on, the unit must follow specific start-up curves that depend, among other things, on how long the unit has been off. Many similar constraints exist, and we will not enumerate them here. Thermal plants can be suitably modelled such that the resulting optimisation problem would be tractable if there were only one unit to meet the demand - but this is certainly not the case with 150.

A hydro-valley is a set of interconnected reservoirs and associated power plants, each of which has a certain number of turbines. Some plants are equipped to pump water from their downstream reservoir into their upstream reservoir (thus allowing this water to be reused later during peak hours), but a plant cannot simultaneously discharge (producing energy) and pump (consuming energy). Using convenient simplifying assumptions, hydro-valleys



Source: www.structure.net. Photo by Dr. Klaus Janberg

Barrage de Chaudanne, a hydro-valley operated by EDF, the French electricity board.

can be modelled as an ordinary linear programming problem, which is once again easy to solve with one valley, but not with fifty.

One way of tackling the unit-commitment problem is price decomposition: the constraint of meeting demand is replaced by a 'dummy benefit' rewarding each megawatt produced by a unit (the marginal benefit is the same for all units). This replacement, called Lagrangian relaxation, allows a separate optimisation of the 150+50 production units. As a result, we are faced with a coordination problem, which is to find optimal dummy prices. Since October 2002, EdF has been using the so-called bundle method, which was invented at INRIA in the mid-seventies, and which is particularly efficient for this sort of problem.

The question then arises: can it be guaranteed that coordination produces equilibrium? The answer is a definite 'no' -

the schedules computed by Lagrangian relaxation are far from meeting demand within acceptable tolerances. While a certain combination of them does meet it, this does not produce a technically feasible schedule.

To balance these two opposite outputs, a heuristic technique has been developed, also at INRIA, and is now being industrially evaluated by EdF. Experiments on real datasets indicate that substantial gains are obtained, as compared to the implementation currently in operation at EdF. In fact, the coordination phase delivers not only optimal dummy prices, but also a lower bound on the minimum possible cost of any feasible schedule; this is a well-known property of Lagrangian relaxation. Usually, the feasible schedules produced by the current implementation cost about 1% more than this lower bound. The new approach typically cuts this excess by a half.

However, the main advantage of the new approach is its robustness, which can be illustrated as follows. On a particular day in 2002, an uncommon peak load of 70GW occurred. The schedule obtained by the current implementation produced an intolerable mismatch of 1.3GW; in contrast, the new heuristic technique managed to find a schedule that cost less and, more importantly, met the demand to within 50MW - quite a reasonable tolerance.

Please contact:

Claude Lemaréchal, INRIA

Tel: +33 04 76 61 52 02

E-mail: Claude.Lemarechal@inria.fr

Articles in this Section

- 45 Solving Complex Problems with Advanced Techniques**
by Enrique Alba, University of Malaga/SpaRCIM
- 47 Nipper - the Neat Internet Protocol Packet Editor**
Sami Pönkänen and Marko Lyijynen, VTT
- 48 On Describing B2B Processes with Semantic Web Technologies**
by Santtu Toivonen, Tapio Pitkäranta, VTT and Jung Ung Min, Stanford University, USA
- 50 Tracking Assets and Vulnerabilities in Corporate Networks**
by Dieter Gantenbein, IBM Zurich Research Lab/SARIT
- 51 Numerical Simulation through Web Services**
by Stephan Springstubbe, Jürgen Klein and Ottmar Krämer-Fuhrmann, Fraunhofer SCAI
- 53 Advanced Communication Networks**
by Michel Mandjes, CWI
- 54 Information Security Research at NTNU**
by Stig F. Mjølsnes and Ingvild Ytrehus, NTNU
- 55 Supervised Term Weighting for Automated Text Categorization**
by Franca Debole and Fabrizio Sebastiani, ISTI-CNR
- 56 Remote Monitoring of Health Conditions**
by Dirk Husemann, IBM Zurich Research Lab/SARIT

Solving Complex Problems with Advanced Techniques

by Enrique Alba

The area of complex problems and modern solving techniques is one of the multiple research lines in progress at the University of Málaga. It has been addressed in the research projects 'MALLBA' and 'TRACER'.

Complex problems – those that can only be solved in non-polynomial time – are becoming common in all the domains of our lives: telecommunications, economics, bio-informatics, industrial environments etc. In these and other fields of research it is often essential to model and solve optimisation, learning, or searching tasks for applications that do not admit an easy formulation. In fact, it is often the case that the problem at hand is non-differentiable, has a large number of constraints or objectives, does not admit contour conditions, or is ill-defined. Also, most real-life problems show a high degree of inter-parameter linkage (epistasis), many locally optimal solutions (multimodality), and present a high dimension.

All these ideas were included in the recently completed MALLBA project (2000-2003), and are being extended in TRACER (2003-2005), a new set of proposals for dealing with complex problems. The MALLBA project was an effort to develop an integrated library of skeletons for optimisation containing exact, heuristic, and hybrid techniques. One important outcome of the project was the MALLBA library, which is publicly available at <http://neo.lcc.uma.es/mallba/easy-mallba>. The three target environments considered for solving complex problems are sequential, LAN, and WAN computer platforms. The list of modern techniques we consider in our studies with MALLBA and TRACER includes evolutionary algorithms, simulated annealing, memetic algorithms, ant-colony algorithms, scatter search, and a set of variations on well-known techniques like branch-and-bound, randomised local search, etc.

Software skeletons are somewhat similar to simplified software patterns. In

MALLBA, software skeletons for all these algorithms have been built with a common internal and public interface. This permits fast prototyping and transparent access to parallel platforms. MALLBA skeletons distinguish between the concrete problem to be solved and the solver technique. Skeletons are generic templates to be instantiated by the user with the features of the problem. Much work has been put into designing reliable, general algorithms, and a permanent interest exists in using software engineering concepts for

parameters). On the other hand, required classes specify information related to the problem. Each skeleton includes the 'Problem' and 'Solution' required classes that encapsulate the problem-dependent entities needed by the solver method. Depending on the skeleton, other classes may be required.

Therefore, the user of a MALLBA skeleton need only implement the particular features related to the problem. This speeds up considerably the creation of new algorithms with minimum effort,

especially if they are built up as combinations of existing skeletons (hybrids). Figure 1 shows an example of design for a simulated annealing algorithm.

The focus of MALLBA was mainly in creating the software and hardware infrastructure (especially an operative WAN) for later developments, while accounting for parallelism and hybridisation.

Many application domains have been addressed within MALLBA, but at the same time, many important issues concerning the problems and the techniques themselves were deferred for TRACER.

In the new TRACER project (<http://tracer.lcc.uma.es>) therefore, powerful and robust algorithms are being developed in close connection with the Internet. Again, in all our research we pay close attention to efficiency and to the software design aspects of the algorithms. This project represents

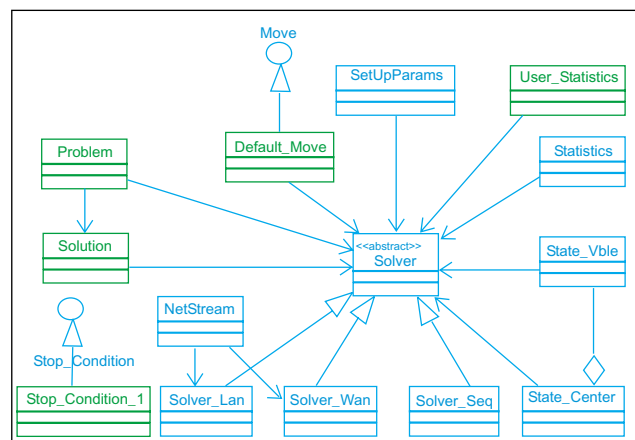


Figure 1: UML design of a simulated annealing heuristic in MALLBA.

correct design. All the knowledge related to the solver method (eg parallel considerations) and its interactions with the problem are implemented by the skeleton and offered to the user.

Skeletons are implemented by a set of required and provided C++ classes that represent an abstraction of the entities participating in the solver method. Provided classes implement internal aspects of the skeleton in a problem-independent way. The most important provided classes are 'Solver' (the algorithm) and 'SetUpParams' (configuration

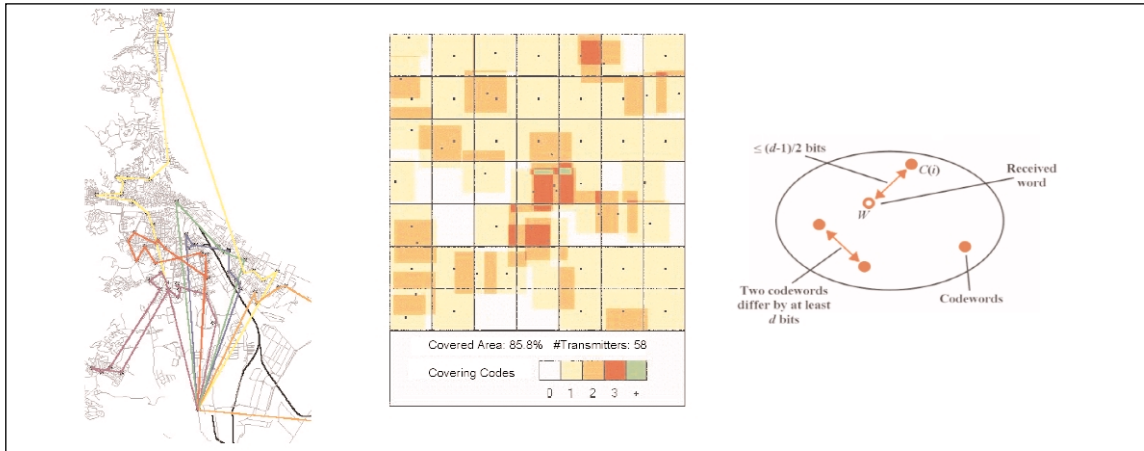


Figure 2: Example of a real VRP in Málaga, location of antennae (left, centre) and error-correcting codes.

a joint effort of five groups in Spain, coordinated from the University of Málaga and devoted to proposing solutions and algorithms (in C++ and also in Java) that outperform existing cutting-edge optimisation techniques. We are working on two fronts: achieving results having the best-so-far accuracy for the considered problems, and developing algorithms for these that show the best-so-far numerical efficiency. At the end of the project we expect to provide the research community with leading algorithms and accurate results, thereby enhancing scientific knowledge on optimisation. Some successful case studies include the following:

- Computing optimal routes for the vehicle-routing problem (vehicular technology). VRP is an important NP-hard problem for which we are able to design minimum cost routes for a fleet of vehicles with a cellular evolutionary algorithm. In this algorithm, tentative solutions evolve in overlapped neighbourhoods connected in a 2-D grid topology. Using local searching in each iteration (specifically 2-opt plus λ -interchange algorithms) has been the key step in outperforming other algorithms in wall-clock time and accuracy. Real-life instances taken from the city of Málaga have been solved to optimality quite efficiently (see Figure 2 left). In Figure 2 (centre) we show a partial solution covering 85.6% of the area. The colour codes indicate whether a given location is covered by 0, 1, 2, 3 or 4 antennae. An optimum solution will have all the area (100%) covered by exactly one single antenna (colour code 1).
- Designing optimal error-correcting codes (information theory) showing the highest minimum Hamming

distance between codewords (see an example in Figure 2, right). For this problem, distributed evolutionary algorithms and a new local search algorithm inspired by the repulsion of charged particles have been used. The numerical effort of our algorithms makes it possible for the first time to obtain solutions for large instances of this problem, in which at present the optimum can only be said to belong to a certain range of values.

- The optimal location of antennae (telecommunications) is another hot topic of research, since UMTS, mobile, and wireless platforms need a new set of networks for their development over the whole of Europe. We have proposed distributed algorithms able to compute solutions covering 100% of the area with a minimum set of strategically located antennae. While in the past this problem required around forty machines working for many hours to reach a solution, our algorithms can now provide us with optimum solutions in a few minutes using between one and eight CPUs. In fact, we are extending these algorithms to other hard problems suggested by well-known telecommunication companies in Europe, in which a physical model of wave propagation plus detailed information of the geographical area is included.
- Optimisation in continuous domains is an additional line in our research. This kind of problem frequently requires very different algorithms to those found in combinatorial optimisation. We have recently proposed a parallel hypercubic evolutionary algorithm dealing with heterogeneous operations including concepts from fuzzy logic to solve continuous problems: systems of

linear equations, computing filters for sound waves, polynomial fitting (Chebyshev), and training neural networks among others.

- The new and exciting field of grid computing is also being addressed to solve some of the most difficult problems. Condor and Globus are two of the most successful technologies we have been using, just as XML and SOAP have been studied for some cases. Concretely, we have developed a new grid algorithm to compute the exact Pareto front of several hard multi-objective problems by using Condor in a grid with more than 100 computers in our department (<http://neo.lcc.uma.es/Software/ESaM>). At present, we are extending the grid to other sites in Spain and Europe.

We are also working on other topics like bio-informatics or dynamical optimisation (DOP). However, besides working on the algorithmic aspect, we are aware of communities of users demanding solutions for their own problems. This is why we plan to release a public Internet service that will help researchers to solve complex problems by feeding them (written in Java, C++, and other languages) into a client/server system and then selecting one solver from an assorted set of algorithms to obtain a solution with a short round-trip time and low effort. Any kind of comment, idea, or collaboration is welcome.

Link:
<http://neo.lcc.uma.es>
Please contact:
 Enrique Alba, University of Malaga/SpaRCIM
 Tel: +34 952 132803
 E-mail: eat@lcc.uma.es

Nipper – the Neat Internet Protocol Packet Editor

by Sami Pönkänen and Marko Lyijynen

Nipper – the Neat Internet Protocol Packet Editor – is a tool for testing networks and network applications. More specifically it is used to create arbitrary packets and to send them to network or directly to applications. Nipper is written in Java and uses portions of C-code and Linux raw sockets.

The need for a tool like Nipper came up in a networking project when we started testing and debugging our software. Numerous tools exist for network testing; those we use most frequently include tcpdump and ethereal for sniffing, ping for connectivity checking and nmap and nessus for scanning of security holes. However, we were unable to find an elegant, simple tool that could create arbitrary IP packets and send them to network or directly to applications.

Instead of coding a quick and dirty, use-and-throw-away type program, we decided to design and implement a customisable tool that would have a simple, user-friendly GUI. In the first place, it would be used quite extensively in the project, and therefore needed to be user-friendly. Further, we also wanted to be able to use it in other projects, meaning it needed to be customisable and extensible, and if possible, able to run on OSes other than Linux (although Linux is what we use most frequently). Java was therefore the natural language to choose.

Nipper turned out to be very usable, and we thought it might prove popular as a starting point for building other tools. Consequently we decided to go open source with Nipper. Nipper was released under VTT Public License.

Nipper's Architecture and Features

Nipper has three main components: PacketEditor is the GUI, PacketEngine is where all the real work is done and Network Interfaces are the components that deliver the packets. See Figure 1 for a diagram of the Nipper internals. The reason for a three-layer architecture is customisability. We wanted the ability to switch between Network Interfaces, to add new protocols to PacketEngine and maybe someday to use one GUI to

control many PacketEngines on different hosts.

PacketEditor takes its input from the user. The user can create new packets, add protocol headers and payload to them and fill in values for different header fields. This input is stored in an XML format description that can be saved and loaded from a file. An example of an XML description is shown in Figure 2.

PacketEditor passes the XML description to PacketEngine where the real action happens. PacketEngine builds a number of real packets from the XML input. Inside PacketEngine, different protocols are represented by Java classes

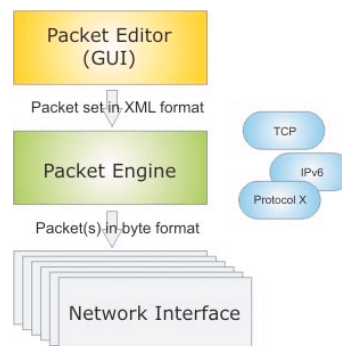


Figure 1: Nipper architecture.

```
<?xml version = '1.0' encoding = 'UTF-8' ?>
<packetset>
  <packet stime="1">
    <header type="IPv4">
      <srcaddr>10.0.0.1</srcaddr>
      <dstaddr>10.0.0.2</dstaddr>
    </header>
    <header type="TCP">
      <dstport>80</dstport>
    </header>
    <header type="Payload">
      HTTP/1.0<data>GET http://www.vtt.fi
    </header>
  </packet>
</packetset>
```

Figure 2:
XML description of a HTTP request.

that perform the actual conversion of XML parameters to byte-format protocol headers. Each protocol class has a set of rational default values for fields which the user has left blank. Tasks like checksumming and calculation of length fields are also done within the protocol classes. Nipper currently includes basic support for IPv4, IPv6, ICMPv4, TCP and UDP protocols and payload data. Partial implementation of Mobile IPv4 is also included.

After the packets are built, PacketEngine starts sending them one by one. This happens by passing the byte array to a Network Interface component. At present there are two Network Interfaces to choose from: PacketDumper and PacketSender. As the name implies, PacketDumper simply dumps the byte-format packets to a debug window. PacketSender uses the Java Native Interface (JNI) and Linux raw sockets to send the packets to network. Currently PacketSender is able to send both IPv4 and IPv6 packets.

Using Nipper

So let's take Nipper for a test drive. As an example we show you how to build and send a packet with IPv4 and UDP headers and a string 'hello world!' as payload. Note that using raw sockets (the PacketSender) requires root privileges.

Figure 3 shows the main window of Nipper. In the lower part of the window is the set of packets with one packet visible. Select the packet and pick IPv4 from the tabs in the top part of the window. Fill in values for different IPv4 fields or leave some fields blank and let Nipper use default values. In Figure 2 we have set the source and destination addresses and the TTL field. Click on 'Add to packet' to add the IPv4 header to the packet you selected earlier. Do the

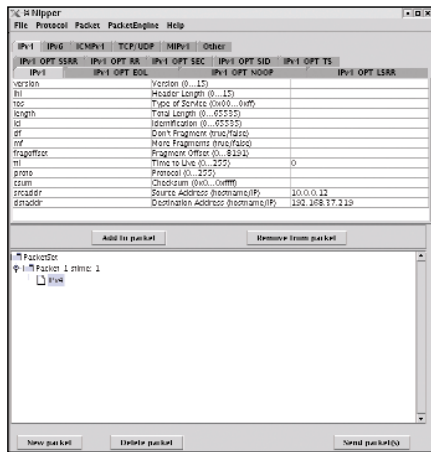


Figure 3: Nipper mainwindow.

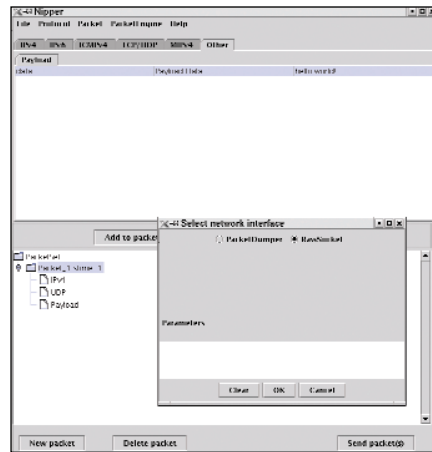


Figure 4: Network Interface selection.

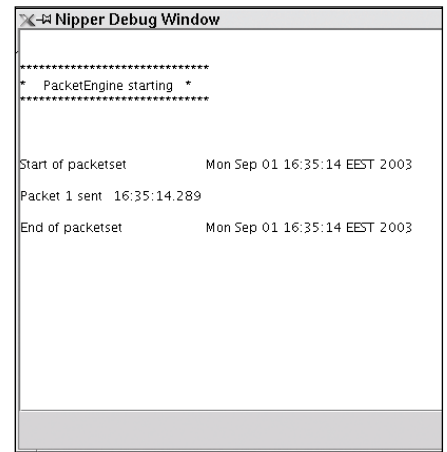


Figure 5: Nipper debug window.

same for UDP and Payload (under the 'Other' tab).

Next you need to select a Network Interface. Go to the 'PacketEngine' menu and 'Set network interface'. Available Network Interfaces are listed on the dialogue box shown in Figure 4. Select 'RawSocket'.

Now you are ready to send the packet to network. Open a shell and start tcpdump if you want to see whether the packet causes any reaction. Then click on 'Send packet(s)'. A debug window opens up and if everything goes as planned you

should see some activity in your tcpdump and a report as shown in Figure 5.

The verbosity of debug messages can be adjusted via the 'PacketEngine' menu. For more information about using Nipper, consult the Nipper User's Guide, which is available at the open source site of VTT.

Developing Nipper

Adding new protocols to Nipper is quite straightforward. We suggest taking a closer look at one of the existing protocols and reading the files under doc/ in the Nipper package. The source code includes JavaDoc documentation.

Network Interfaces can also easily be added. One could for instance create a Network Interface that uses the tun/tap device instead of raw sockets. Again, take a look at existing source code and the documentation bundled with Nipper!

Link:
<http://opensource.erve.vtt.fi>

Please contact:
 Marko Lyijynen, VTT
 Tel: +358 9 456 4311
 E-mail: marko.lyijynen@vtt.fi

Sami Pönkänen, VTT
 Tel: +358 9 456 7232
 E-mail: sami.ponkanen@vtt.fi

On Describing B2B Processes with Semantic Web Technologies

by Santtu Toivonen, Tapio Pitkäranta and Jung Ung Min

The PAX project, carried out by VTT Information Technology in cooperation with the Center for Integrated Facility Engineering (CIFE) at Stanford University, concentrates on Intelligent Web Services within the context of two specific B2B cases, namely service chain composition and expediting process.

Two emerging trends are reshaping the Web as we know it today. One of these is the Semantic Web, in which content is consumed by computer programs in addition to humans. The other is Web Services, which provides a set of standards and practices for the remote use of computer applications. The combination of these two is sometimes referred to as Intelligent Web Services. The PAX project concentrates on this within the context of two specific B2B cases, namely service chain composition and

expediting process. More specifically, the cases are approached in a 'bottom-up' manner, meaning that service providers have a degree of freedom in describing and executing their services. This in turn means that some level of intelligence or adaptability is expected from the service consumers, that is, the service composer and the expeditor.

Web Services carry on the tradition of distributed computing. Individual services are scattered over the Internet

and are used remotely. Web Services are modular and composite by nature, meaning that simpler services can be grouped together to form more complex ones. It is necessary that the service composer be able both to find the services and also to evaluate them before the composition. Service discovery is easier to achieve, since Web Services are addressable with URIs, which can be stored in a common repository such as UDDI. Service evaluation is not so straightforward, because it presupposes

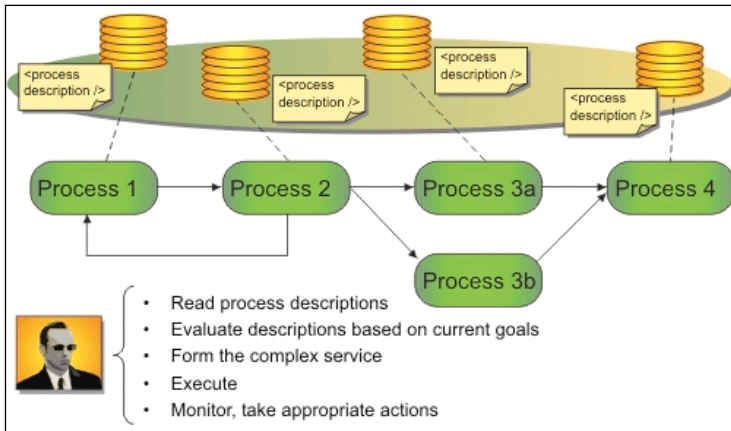


Figure 1:
Composing and executing a complex Web Service.

descriptions of the services in a commonly agreed-upon fashion. This problem is especially evident if the service composer is a computer program. Service descriptions must be serialised in a machine-understandable format, and conform to a common ontology.

The richness of service descriptions can vary from plain service-type categorisation to detailed descriptions about service characteristics. The PAX project concentrates especially on process-related service descriptions. Services are described using Semantic Web technologies such as RDF and OWL. Figure 1 depicts a complex Web Service, which is composed based on the process descriptions of the individual services.

In the PAX project, the service composition task is approached in terms of a specific scenario from the construction industry. In this scenario there is a subcontractor (SC), sub-subcontractors (SSCs), banks, and material suppliers. The scenario also contains three kinds of

processes: informational (such as requests for quotes), material (such as product deliveries), and financial (money flow). Some of the processes are one-to-one (eg product deliveries) whereas others are one-to-many (eg biddings for SSCs and banks). The scenario can also be configured, based for example on whether the project is on or behind schedule. For instance, the service composer, in this case the SC, can choose the more expensive but faster SSC if the project is behind schedule – something that it would not do if the project was running on time.

Apart from this type of service composition task, the adaptability of software agents is also investigated. More specifically, agents adapt to interaction protocol descriptions if unaware of the utilised protocols beforehand. Interaction protocols are typically used when modelling software agent conversations. They too are described using Semantic Web technologies, and conform to a common interaction protocol ontology.

The specific case chosen for studying interaction protocol-based adaptability is an expediting process, depicted in Figure 2. The general contractor (GC) expects status reports from SCs regularly. Suppose it has received one from some SC, containing information about an important steel delivery that has not yet arrived. GC designates an expeditor to contact other process participants directly in order to verify the status of the steel delivery. Traditionally this means paying visits to the shop floors or making phone calls. Instead, expediting is now approached with novel technologies, namely software agents representing the process participants and Semantic Web languages for describing the utilised interaction protocols.

From SC's status report, GC is able to extract the steel order ID as well as contact information for the steel supplier (SS) and an external carrier (CR). With the order ID, it asks SS for the delivery ID of that order. Note that GC is initially unaware of the interaction protocols provided by SS and CR. Rather, it is aware of a common interaction protocol ontology according to which the individual protocol descriptions are designed. With that knowledge, it is able to adapt its behaviour to correspond to SS's interaction protocol.

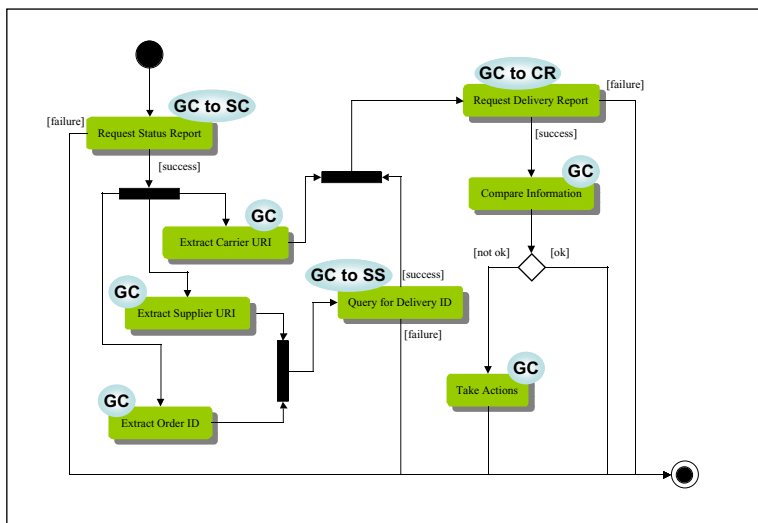
With the delivery ID provided by SS, GC then contacts CR and requests a delivery report. Like SS, CR also offers descriptions of the interaction protocols it uses to communicate, and GC then adapts its behaviour to conform to those. After receiving the delivery report, GC can compare the information in the status report and the delivery report and take action if needed. The comparison and possible resulting actions can be performed either by the software agent itself or, more realistically, by a human being.

Link:
<http://www.vtt.fi/tte/proj/pax/>

Please contact:
Santtu Toivonen, Tapio Pitkäranta VTT
E-mail: santtu.toivonen@vtt.fi

Jung Ung Min, Stanford University, USA
E-mail: jumin@stanford.edu

Figure 2:
Interaction protocol-based expediting process.



Tracking Assets and Vulnerabilities in Corporate Networks

by Dieter Gantenbein

In today's dynamic information society, organizations critically depend on the underlying computing infrastructure. Mobile users and computing devices add to the challenge. Business operations, intellectual property, and corporate value are quickly at risk. The IBM Zurich Research Laboratory has developed a tool called Intelligent Device Discovery (IDD) that collects information from all possible sources, computes an aggregate picture of assets, and categorizes their usage and vulnerabilities.

Tracking computing devices as assets with usage, health, and vulnerability information facilitates the provision and maintenance of an efficient, optimized service. Recent incidents with viruses and worms appear to support this. In general, a precise understanding of the operational infrastructure is a key element of many corporate decisions. Examples also include the negotiation of outsourcing contracts, the planning of mergers and acquisitions, server consolidation and business optimization.

Building an accurate inventory of computing assets in heterogeneous dynamic systems and networking environments is difficult, especially when only limited privileges are available and no prior device instrumentation has taken place.


Classical methods for inventory and asset management quickly reach their limit in today's dynamic environments: Periodic physical inventories ('wall-to-wall') have the distinct advantage of identifying the actual location of devices but require costly human visits ('sneaker net') and cannot detect mobile, currently out-of-office equipment nor the existence and use of logical assets. Financial asset tracking, although an accepted process in its own right, cannot detect additional equipment brought into or accessing the resources of an organization. Periodic self-assessment questionnaires to be filled out by individual end users or their cost-center managers are another, often complementary approach. Apart from the human effort they require and the inaccurate, incomplete data they may contain, most forms pose questions the answers to

which could easily be obtained from the infrastructure itself.

Well-managed computing infrastructures typically equip servers and end-user devices with software daemons to track the system's resources and health. There are

with? How can you quickly substitute in any missing data?

How many old Windows systems with less than 512 MB of memory do we need to upgrade this geography? Where does this device with a strange MAC address



Networks, communication paths, and sources of information.

many situations, however, in which these daemons cannot be relied upon. In many organizations, there are a fair number of devices that are brought in ad-hoc without appropriate instrumentation, for which instrumentation is not available, or on which instrumentation has been disabled.

It therefore appears advisable to complement process and policy-based asset and security management with 'automatic sensors' to recalibrate the dynamic and heterogeneous environment. This is where IDD fits in.

Scenarios

Your organization may already have detailed inventories, but what about the IT infrastructure of the company you interact

come from? Which machines are unsafe, ie, without the latest patches and not protected by antivirus and security software?

Computer A is a safe machine used by a secretarial employee on the mornings of four out of seven days a week. Computer B — perhaps a student-lab workstation — has many server ports open, is not completely safe, and is shared by users X, Y and Z. Some server has local unprotected databases and no screen password set. A portable device currently in office C has no hardware password set.

And those rogue servers should either be turned off or turned into a policy-based operation. We may want to send e-mails to those employees whose machine

configurations harbor discovered or accepted risks. Discovery provides the means to detect early potential exposures and track down the exact location of assets in violation of current policies.

Information Sources

Corporate networks of large companies span geographies, the individual locations of which typically consist of separate segments for user access, servers, as well as intranet and extranet connection zones. The operation of wide-area and/or campus networking, servers, and/or user devices may be handled by different organizations. In general, there are neither universal administrators nor "one size fits all" credentials. In order, nevertheless, to be able to derive an overall picture, we propose that data be harvested from whatever source possible, and then combined into an integrated information model. The range of information sources considered encompasses network-based autodiscovery, privilege-based additional access to online network and application subsystems, and the lean processing of manually provided ledger tables containing financial and physical inventories and configurations. Bothering the end user is considered only as a last resort.

Techniques

The sources of information are manifold, and the techniques to access them are even more so. The communication stack and the device ports can be actively mapped by scanning from administrative machines, including IP ping sweeps, UDP/TCP port scans, and remote Windows SMB/Registry/WMI fingerprinting. Should established policies limit the yield of this technique for some of the target devices, analysis of subnet traffic can also passively map otherwise stealthy devices. Although this approach assumes probes on appropriate segments and VLANs, the trend to corporate netflow architectures that collect aggregated traffic accounting records from strategic subnets provides for a more strategic access-point for network flow information. Observing the communication between devices can be also used to log networking services, including (DNS) name servers containing directory information about registered resources, DHCP address servers with lease information, WLAN wireless access servers, and PIX firewalls. Some of the related events may also be readily available at syslog servers. Whereas managed end-user systems may be instrumented with agents that interact

periodically with corporate management servers, eg to keep software up-to-date or to download new virus and worm defense policies, there is also the entire realm of other application servers for mail, group communication, Web applications, mobile-user access portals, etc., that may provide detailed information about usage patterns.

Conclusions

IDD is a network-based IT asset discovery and categorization tool. It combines various device discovery, network and security scanning programs with enhanced data collection, mining techniques, and distributed automation to form a single application built upon the IBM e-business platform with WebSphere and DB2 on Windows and Linux. It is used today by consulting, outsourcing, and security teams on internal and customer networks.

Link:

<http://www.zurich.ibm.com/idd>

Please contact:

Dieter Gantenbein,
IBM Zurich Research Lab/SARIT
Tel: +41 1 724 88 53
E-mail: dga@zurich.ibm.com

Numerical Simulation through Web Services

by Stephan Springstube, Jürgen Klein and Ottmar Krämer-Fuhrmann

Simulation portals provide secure access to data and the use of different analysis and simulation tools via Internet. Benefits are a safe access independent of location or time, personalized services with an adequate user support, and modern security methods like certificate authentication or encryption. Thus simulation portals meet industrial requirements in a 'portable' environment.

The expansion of the Internet and the development of secure access and communication methods enable a new and integrative usage of simulation tools. Web portals offer the integration of data, applications, services and resources in a personalized environment.

The Internet or Intranet enable developers to access central resources such as databases or computer systems via a graphical user interface. This does not only reduce investment costs, but also

keeps expenditures for installations and support as well as licenses at a minimum.

From Authentication to Simulation

Important components of a web portal are:

- user management – the basis for authentication
- object management – for the realization of 'user workspaces'
- control of processes incl. accounting of computer resources and license fees.

Modern authentication methods, safe communication and quick database connections allow clients to use resources via the portal independently of time or location.

Numerical Simulation in a Web Portal

Through increased computer power numerical simulation is becoming more important for the development of turbo machines. The software TRACE (Turbo-machinery Research Aerodynamic

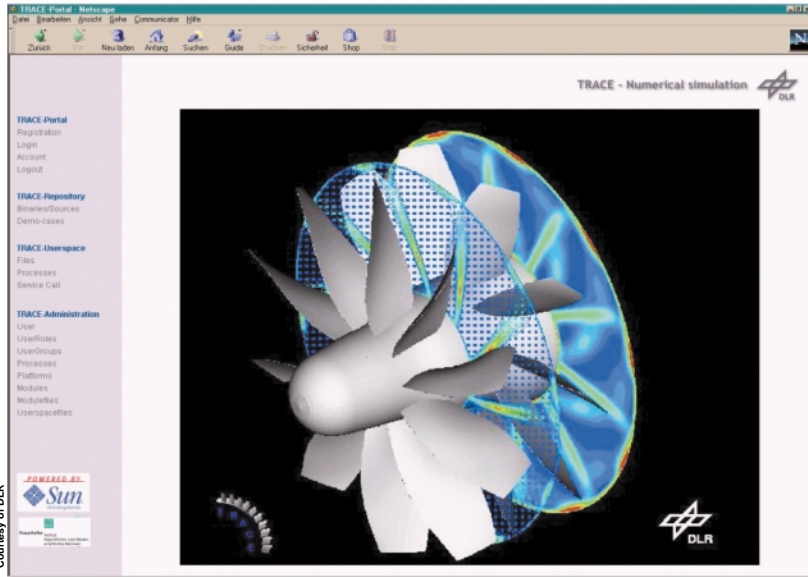


Figure 1: Time-accurate 3D-Simulation of a Propfan with TRACE.

Computational Environment), developed by the Institute for Propulsion Technology of the German Aerospace Center (DLR) in Köln-Porz, provides a detailed insight into complex fluid processes.

Fraunhofer SCAI integrated this software solution into a web portal, thus enabling a better use by the industry.

Individual Services for Users – the TRACE-Portal

Web portals cannot only be used easily by means of a browser, but also offer a personalized Internet environment and user guidance. This supports perfectly a variety of user roles like Internet surfers, registered users, licensed customers and administrators.

In particular the role und group definition of users makes it easy to support various development teams in the best possible way. The portal solution is portable to other computer systems and can be customized to every application and service. Thus portals can easily be adapted to new technologies.

Client-Server Technology on the Basis of Sun ONE™

The software Sun Open Network Environment (Sun ONE™) is the basis for the implementation of the TRACE Portal. Sun ONE™ is Sun's standards-based software platform for building Web Services on Demand. Originally being developed for e-commerce applications, now Sun

ONE™ has been adapted by Fraunhofer SCAI for scientific applications. Sun ONE™ supports web portal solutions with a modular architecture. The TRACE Portal consists of three components (see Figure 2): Web container, Enterprise Java Beans container and Compute server. Each of these components realize a web service interface and is therefore easily adaptable to new services or technologies. In the following we describe the portal architecture.

Client

The simplest access to the portal server is via a browser. The browser realizes

access to the public area by simple HTTP protocol, and for registered users to the restricted area by the HTTPS protocol using ssl encryption.

Alternatively web services or applications can access the simulation service directly. In both cases the client system communicates via Internet connection with the web portal server.

Web and Application Server

Server Module 'Web Container': The server module accessed by the client is the web container. It realizes the presentation layer forming the user interface. It is implemented by Servlets and Java server pages (JSP). The web container executes the user authentication and guides through all portal services by means of a graphical user interface. Different user roles are supported by the portal, so that each user works in a personalized web environment with individual services, account and workspace data.

Server Module 'EJB Container': The Enterprise Java Beans (EJB) container realizes the business logic and can be understood as an interface between the user and resources, like the database or the compute server. The EJB container is realized in Java.

Server Module 'Database': The database stores all data of the portal, like user

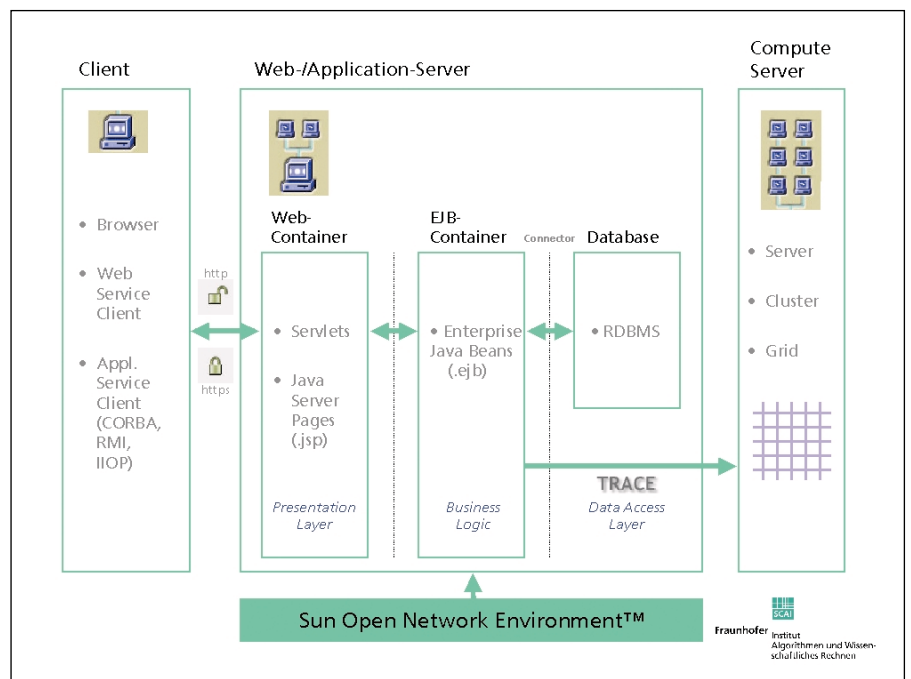


Figure 2: Architecture of the TRACE-Portal.

data, virtual user workspace, version management and account data.

Compute Server

The most important web service for the user is the execution and control of numerical simulations. This is controlled by Enterprise Java Beans of the EJB container. The data access layer allows to upload data to the computing platform. TRACE simulations are executed on a multi-processor Sun-server or on a computational GRID.

The user can monitor and control computations. He is able to configure, start, stop or terminate simulation processes. Accounting is done automatically by the portal.

The portal is well suited for the use, if several departments are working with the same simulation code. Portal solutions support access to central resources, which minimize license fees, man power of installation, updating, and administration efforts.

Fraunhofer SCAI offers complete solutions for the conception, design and realisation of innovative Internet presentations. Portal solutions offering numerical simulation as a web service in a safe, user-friendly web environment, which can be used independent of location or time.

Links:

<http://www.scai.fraunhofer.de/web>
<http://www.trace-portal.de>

Please contact:

Ottmar Krämer-Fuhrmann
Fraunhofer SCAI
E-mail: ok@scai.fraunhofer.de

Advanced Communication Networks

by Michel Mandjes

The current trend in telecommunications is towards ubiquitous availability of any application - at any time or any place. This will eventually lead to the universal accessibility of a common network infrastructure that supports all imaginable communication services. CWI's researchers are working on the use of queueing theory in communication networking, which is relevant in many areas, including wireless services and network economics research. Application of the developed techniques could lead to a more economic use of network resources.

Technological innovations in networking are occurring at a rapid pace. Whereas early networks supported a limited number of communication services, nowadays a broad variety of services are integrated into a single infrastructure. This inevitably leads to performance issues. For instance, a voice user asks for essentially real-time service, in particular minimal delay, whereas these delay requirements are far less demanding when sending an e-mail. While the network could be designed to meet the most stringent of these Quality of Service (QoS) requirements, this

would clearly lead to low utilisation, and consequently a waste of resources.

CWI's Advanced Communication Networks theme (PNA2) aims to apply queueing theory to tackle these traffic management problems. Queueing theory is all about congestion effects arising from scarce resources. It helps to understand how congestion arises, but also how congestion can be avoided. In particular, queueing theory facilitates the development of QoS differentiation mechanisms. CWI's research has focused on novel scheduling mechanisms, such as Generalised Processor

Sharing (GPS), that are capable of offering multiple QoS levels. A key research issue in this area concerns the protection of user classes against 'misbehaving' traffic streams – in this way GPS mitigates the impact of heavy tails. This strand of research requires sophisticated probabilistic techniques such as large-deviations methods, as well as a thorough knowledge of Gaussian processes.

Recent developments in networking have led to challenging new research directions. One important field concerns performance support for wireless services. Current wireless systems like GSM are highly customised for carrying voice traffic, and offer only limited low-bandwidth data applications, such as short messaging services (SMS). The next generation of wireless networks like UMTS are specifically designed to support a wide range of high-speed data applications, such as Web-browsing sessions, document transfers, and imaging services, in addition to conventional voice calls. The integration of these heterogeneous applications raises similar challenges to those described above for wireline systems. However in



Due to both the ongoing increase in traffic volume and the introduction of new network services, major infrastructural upgrades are necessary, such as here in the Sciencepark Amsterdam. The networking group at CWI explores the trade-off between efficient use of network resources and the quality of service offered.

Photo: CWI

wireless environments, these issues are further exacerbated by the relatively low bandwidth and scarcity of spectrum (see the UMTS licence auctions), and the propagation characteristics of wireless signals.

In particular, the transmission rate in wireless communications fluctuates considerably over time due to multi-path propagation effects. These rate fluctuations provide an opportunity to achieve throughput gains by scheduling data transmissions. Recent research in PNA2 has explored the QoS as perceived by users of such 'channel-aware' or 'opportunistic' scheduling algorithms. It was shown that user-level performance may be evaluated by means of a processor-sharing model with a variable service rate. The variable service rate captures

the throughput gains achieved by the channel-aware scheduling algorithms.

A second major effort has been invested into network economics research. As argued above, a prerequisite for the support of heterogeneous applications is the ability to offer multiple performance levels. A next step is to decide how these performance levels should be priced. Evidently, with no pricing, or with a simple flat fee, all users would opt for the premium class, disabling performance differentiation. This can be seen in today's Internet: the flat fee does not encourage 'economical use' of resources. The consequence is that the Internet tends to be used by those who are relatively insensitive to congestion, thus pushing away services that require stringent performance guarantees (such as voice users). This cross-disciplinary

phenomenon is usually called the 'tragedy of the commons'. A way to prevent this 'tragedy' from happening is by instituting charges. Performance-sensitive users are usually willing to pay more for the service than congestion-indifferent users. Hence by choosing the prices of the different performance levels appropriately, incentives can be given such that only the performance-sensitive users join the premium class. Recent research in PNA2 has focused on this type of problem, with our approach incorporating microeconomic and game-theoretic elements into the queueing framework.

Link:

<http://www.cwi.nl/pna2>

Please contact:

Michel Mandjes, CWI

E-mail: Michel.Mandjes@cwi.nl

Information Security Research at NTNU

by Stig F. Mjølsetnes and Ingvild Ytrehus

NTNU, the Norwegian University of Science and Technology has established an interfaculty Information Security Research Program, organised under the university's strategic focus area of Information and Communication Technologies.

This graduate education and research cooperation initiative is currently run by the five departments of Telematics, Computer and Information Science, Mathematical Sciences, Physical Electronics, and Industrial Economics and Technology Management. The overall aim is to address problems of ICT vulnerability at the strategic level by supervising and carrying out research with PhD students on a wide range of topics essential to information security technology. In this way we will train experts with the ability to perform effective information security analyses and provide solutions to ICT vulnerabilities, wherever they occur.

Information Security

Information security is vital for the acceptance and common use of information systems in most sectors of activity in our information-intensive society, including health, finance, trade, public and private administration, media and entertainment, communication and transport, and within

the ICT industry itself. Open networked systems are particularly challenging in this regard. Clearly, Internet services and mobile networks already suffer from this kind of problem, resulting in a great demand for comprehensive research and the training of experts within information security technology.

Currently, technology and methods for information security involve:

- hardware-based mechanisms (tamper-resistance, OS kernel support, and signal processing)
- cryptology (mathematical primitives, protocols and models)
- software engineering (language/mechanisms/tools)
- software systems (operating systems, database systems, middleware platforms)
- networked systems (communication protocols, naming, routing, adaptivity)
- methodology (formal logic, evaluation criteria, threat and vulnerability analyses, audits, best operation practice and policies).

The difficulty of analysis and construction in information security technology grows rapidly with increasing scalability, functionality, resource distribution and partitioning of security policy. Information security requirements place conditions on the design of system structure, user interfaces, data storage, processing locality, communication, and management. On the other hand, solutions must meet user expectations and be acceptable within cost restrictions.

Norwegian Research Program

In the financial allocation letter of 2003, the Norwegian Ministry of Trade and Industry requested that the Norwegian Research Council (NFR) initiate a strategic research program within the field of 'ICT Security and Vulnerability' (IKT-SoS). A total budget of NOK 59 million was allocated, and it is to be operational from 2003 to 2007 (1 Euro ^a 8 NOK). Only NOK 15 million of the total budget were allotted for this first year by the council. Nevertheless, total submissions

for research funding amounted to NOK 90 million, indicating a strong interest by Norwegian research groups in working on the varied problems in this field.

NTNU Objectives

The eight-professor committee of the NTNU Information Security Research Program has created and submitted project plans which will contribute directly to the goals set through the National Strategy for ICT Security in general, and by the Norwegian Science Council IKT–SoS. These include the following initiatives, as specified in the call:

- strengthening of national education in information security in both depth and scope
- further strengthening of existing information security research networks in Norway and Europe, and take-up by user organisations
- support for political strategies aimed at reducing ICT vulnerabilities
- pursuit of excellence in information security RTD in Norway.

About fifteen years ago, NTNU started educational activity within this field in

the Department of Telematics. Student interest and motivation are now very encouraging. This year (2003), one of the Masters degree information security courses produced 170 examnants, and about 36 Masters theses in information security were completed in the five cooperating NTNU departments. Next year, the number of PhD candidates is expected to be 12-15, within a framework of ongoing research projects, labs, and international university cooperation.

Research Projects

New research projects that have been defined and are underway are presented as four ellipses in the depicted NTNU reference model, ranging from human-organisational issues ('people') to deep mathematical theories ('machines'). Stage 1 funding by NFR enables commencement of a PhD project in role-based access control, a post-doctoral fellowship in secure protocols, and funding for visiting scientists in all areas indicated. Stage 2 next year will enhance this picture. Additionally, several Marie Curie Fellowships are open for visiting PhD students. Communication security

research forms a significant part of the activity in the recently established Centre of Excellence, 'Quantifiable Quality of Service in Communication Systems' at NTNU, in which around five PhD positions are allocated to security-related studies as of now. Cooperating with NTNU, industrial research projects related to information security are executed at SINTEF and UNINETT, including the National Centre for Information Security in Trondheim.

Links:

Strategic focus area of ICT at NTNU:
<http://www.ntnu.no/satsingsomraader/ikt/>

Information Security Program:
http://www.ntnu.no/satsingsomraader/ikt/fp/info_sikker/index.htm

Q2S: <http://www.ntnu.no/q2s/protrust.php>

Marie Curie Training Site in Reliability, Safety and Security Studies at NTNU:
<http://plato.maskin.ntnu.no/emner/mcts/visartikkel.php?id=18>

NORSIS: <http://www.norsis.no>

Please contact:

Stig Frode Mjøl̄snes, NTNU
Tel: +47 73 55 04 59
E-mail: sfm@item.ntnu.no

Supervised Term Weighting for Automated Text Categorization

by Franca Debole and Fabrizio Sebastiani

Researchers from ISTI-CNR, Pisa, aim at producing better text classification methods through the use of supervised learning techniques in the generation of the internal representations of the texts.

Text classification (TC) is the activity of automatically building, by means of machine learning (ML) techniques, programs ('text classifiers') capable of 'pigeonholing' natural language texts, ie placing them in categories from a predefined set, according to an analysis of their content. Instances of text classification are topic spotting, spam filtering, genre classification, or automated authorship attribution, depending on the nature and meaning of the categories being considered.

The construction of an automatic text classifier is usually articulated in three phases. The first phase is term selection, in which the most relevant terms for the classification task are identified. The

second phase is term weighting, in which document-dependent weights for the selected terms are computed so as to generate a vectorial representation for each document, in which each term is weighted by its contribution to the extensional semantics of the document. The third phase is classifier learning, in which a learning device generates a classifier from the vectorial representations of the training documents.

This entire process involves an activity of 'supervised learning', ie one in which information on the membership of training documents in specific categories is used. Traditionally, supervised learning only enters into phases 1 and 3;

phase 2, instead, usually relies on techniques borrowed from text search such as tf-idf ('text frequency * inverse document frequency'), a weighting function based on the distribution of the terms within the document and within the collection, where a high value indicates that the word occurs often in the document and does not occur in many other documents. As a consequence, these techniques do not exploit the information provided by training documents since text search does not involve any training documents.

In our current work we propose that learning from training documents should also affect the term weighting phase, ie

that information on the membership of training documents in specific categories be used to determine term weights. We call this idea supervised term weighting (STW). As an example of STW we propose a number of 'supervised variants' of $tf \cdot idf$ weighting, obtained by replacing the idf part of $tf \cdot idf$ with the same function that has previously been used in the term selection phase. The rationale of replacing idf lies in the fact that idf represents a measure of the document-independent value of a term, but as such it is suboptimal in a text classification context, in the sense that it relies on an intuition ("the document-independent value of a term is inversely proportional to the number of documents in which it

occurs") that is valid also in information retrieval tasks in which no training data are available. Feature selection functions rely instead on an intuition ('the document-independent value of a term is directly proportional to how differently the term is distributed in the positive and negative examples of the categories of interest') that refers to the presence of categories, and that is thus specific to tasks in which category data is available.

We have run STW experiments on Reuters-21578, the standard benchmark of text classification research, with three classifier learning methods (Rocchio, kNN, support vector machines), three term selection functions (information

gain, chi-square, and gain ratio), and two policies for addressing term selection and weighting ("local" and "global"). Results show that STW is a powerful notion since, when instantiated with the 'gain ratio' feature selection function, it can bring about improvements as high as 11% in accuracy over a standard $tf \cdot idf$ representation with a support vector machine learner.

Links:

<http://faure.iei.pi.cnr.it/~fabrizio/Publications/SAC03b.pdf>

Please contact:

Fabrizio Sebastiani, ISTI-CNR
Tel: +39 050 3152892
E-mail: fabrizio.sebastiani@isti.cnr.it

Remote Monitoring of Health Conditions

by Dirk Husemann

IBM researchers working with medical device manufacturers and mobile-phone handset manufacturers have created a unique solution to track vital health signs.

This technology, if adopted, could help limit a patient's health risk, increase his or her level of comfort, reduce both the number of check-ups and the time spent in hospital, and enhance the effectiveness of pharmaceutical field tests. The solution includes a so-called 'pill box' and a home-based blood-pressure monitor that resembles a large wristwatch.



Health-monitoring devices such as the blood-pressure cuff (left) and pill box (right) send data to the mobile phone via Bluetooth. The mobile hub software integrated into the mobile phone (centre) forwards the data to a care centre for monitoring (screen in the back) and returns reminders or alarms in an emergency.

These new prototypes build on existing devices available through specialist suppliers and diagnostics companies. The new enhancements make it possible for the devices to transmit secure patient data, free of tampering, through a mobile phone to an Internet portal and so to the doctors and nurses monitoring the patient.

The wristband blood-pressure monitoring device can also check other vital signs such as heart rate, and is activated by simply pressing a button. Blood-pressure readings, for example, are gathered from one or more sensors via a Bluetooth short-range radio connection. Once transmitted, secure access ensures only authorised medical personnel see the patient's data. If an unusual reading comes through, either a reminder can be

sent to the patient to take his or her medication or a new prescription can be made, depending on the doctor's diagnosis.

The 'pill box' helps keep track of the patient's medication by sending a signal to his or her mobile phone every time a pill is removed. If a patient forgets to take medication or is taking too many pills, he or she is sent a reminder via

mobile phone to follow the prescribed doses.

Since vital measurements are readily available and can be monitored on-the-go, it is expected that both these solutions will increase patient mobility and help reduce the need for patients to visit the doctor. Regular and frequent monitoring can also bring peace of mind to people with chronic illnesses or those recovering from an operation. These devices are already of great interest to the pharmaceutical industry, which will be able to conduct field trials of new drugs more accurately.

These solutions are examples of how technology from the emerging field of information-based medicine can be used to deliver medical information in real time to help the medical profession and the health industry improve the quality of patient care.

Link:

<http://www.zurich.ibm.com/csc/mobile/index.html>

Please Contact:

Dirk Husemann, IBM Zurich Research Lab/SARIT
E-mail hud@zurich.ibm.com

FM 2003 – 12th International Symposium on Formal Methods

by Stefania Gnesi

More than 200 scientists from academia and industry attended the 12th International FME Symposium FM 2003, an ERCIM sponsored event, hosted by ISTI-CNR, 8-13 September, 2003, at the CNR research campus, Pisa.

FM03 was organised by Formal Methods Europe, an independent association whose purpose is to promote and support the industrial use of formal methods for the development of computer systems. Formal methods have been controversial throughout their history, and the realisation of their full potential remains, in the eyes of many practitioners, merely a promise. Have they been successful in industry? If so, under which conditions? Has any progress been made in dispelling the scepticism that surrounds them? Are they worth the effort? Which aspects of formal methods have become so well established in industrial practice that they have lost the 'formal method' label?

FM 2003 aimed to answer these questions by seeking contributions not only from the Formal Methods community but also from outsiders and even from sceptics. The organisers were confident that the presentation of a wide spectrum of experiences and the opportunity for the open discussion of contrasting opinions would foster a better and deeper understanding, if not a wider adoption of Formal Methods.

The scientific programme chairs of the symposium were Keijiro Araki (Kyushu University) and Dino Mandrioli (Politecnico di Milano). 144 submissions – mainly of high quality – were received from more than 25 countries throughout the world. Less than one third were accepted, thus maintaining the traditionally high standards of the FM symposia. The symposium also included talks by four invited speakers: Kouichi Kishida (SRA, Japan), Brian Randell (University of Newcastle, UK), Gerard Holzmann (NASA/JPL, USA), Jean-Raymond Abrial (Consultant, France) and two presentations by industrial spon-



Stefania Gnesi, general chair of FM'03 addresses the audience.

sors. Eight tutorials, seven workshops, and tool demonstrations were associated with the main programme.

New to the symposium was the Industrial Day dedicated to Formal Methods and Industry, organized by Dines Bjørner (Technical University of Denmark) and sponsored by FME and CoLogNet. There is a strong commitment within FME towards an increasing involvement of industry in the formal methods community. The Industrial Day thus had two purposes: to focus on the industrial use of formal methods within FM 2003, and to launch a new association, ForTIA, the Formal Techniques Industrial Association. The first Chair of ForTIA is Anthony Hall from Praxis Critical Systems, UK.

Links:

<http://fme03.isti.cnr.it/>
<http://fmeurope.org>
<http://www.eurice.de/colognet>

Please contact:

Stefania Gnesi, ISTI-CNR
 General Chair FM 03
 Tel: +39 050 3152918
 E-mail: stefania.gnesi@isti.cnr.it

ERCIM-Sponsored Events

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

Conferences

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

Workshops and Summer Schools

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

Forthcoming Events Sponsored by ERCIM:

- SOFSEM 2004 – The Conference on Current Trends in Theory and Practice of Informatics, Merin, Czech Republic, 24-30 January 2004
- 9th International Conference on Extending Database Technology (EDBT'04), Heraklion, Crete, Greece, 14-18 March 2004
- 27th Annual International ACM SIGIR Conference, Sheffield, UK, 25-29 July 2004
- MFCS 2004 - 29th International Symposium on Mathematical Foundations of Computer Science, Prague, Czech Republic, 22-27 August 2004

Further information:

<http://www.ercim.org/activity/sponsored.html>

ERCIM Workshop on Soft Computing 2003

by Petr Hájek

The second workshop of the ERCIM working group on soft computing was held from 18-20 October in Brno, Czech Republic. It was held in parallel with the DATAKON conference, a high-profile traditional conference focussed on theoretical and technical backgrounds, best practices and development trends in the deployment of information technology for information systems development, including practical results from industry.

The previous three workshops on Soft Computing (with the 2002 workshop the first to be recognised as an ERCIM workshop) were held in parallel with SOFSEM conferences. Unfortunately this year SOFSEM altered its conference date such that it clashed with an important conference on fuzzy logic, leading

to the workshop aligning itself instead with DATAKON. This proved to be extremely helpful, since local administrative matters were attended to by the DATAKON organisers. This included the opening concert and welcome reception on October 18 (the arrival day). Our thanks are due to Professor Staudek from Masaryk University, the chair of DATAKON, for this arrangement.

The 2003 workshop included four invited talks and nine contributions. Proceedings with full papers appeared as Issue 5 of Volume 13 of the international journal, Neural Network World.

Invited Talks

- C. G. Fermüller: 'Theories of vagueness versus fuzzy logic: can logicians learn from philosophers?'
- B Gerla: 'Many-valued logics and semirings'
- S Jenei and F Montagna: 'A proof of standard completeness for non-commutative monoidal t-norm logic'

- P Jipsen: 'An overview of generalised basic logic algebras'.

Contributed papers from Czech, Slovak, British and Italian authors ranged over various topics of probabilistic, possibilistic and fuzzy approaches to managing uncertainty.

To quote from my editorial in the proceedings: "The fourth workshop on soft computing appears to continue in our good tradition and to contribute to foundational and other aspects of the discipline of soft computing. It is hoped that we shall keep this tradition also in the future."

Link:

<http://www.datakon.cz/datakon03/softcomp.html>

Please contact:

Petr Hájek
Institute of Computer Science, Academy of Sciences of the Czech Republic/CRCIM
E-mail: hajek@cs.cas.cz

ISC 2003 - Industrial Simulation Conference

by Ana Pajares, Juan C. Guerri and Carlos Palau

The first Industrial Simulation Conference (ISC'2003) was held in Valencia, Spain, from 9-11 June 2003. This annual international conference is intended to give a complete overview of recent industrial simulation related research for the European community and the rest of the world.

The meeting, hosted by the Communications Department of the Technical University of Valencia, was organized by the European Simulation Society (Eurosis) and co-sponsored by the Ecole Nationale Supérieure des Arts et Industries Textiles (ENSAIT) and the Ghent University. With the objective of exchanging techniques and ideas among universities and industry, ISC'2003 put

together both scientific research results from universities and practical approaches from industry, which participated with a variety of contributions to the following topics:

- Methodology
- Verification and Analysis
- Complex System Simulation
- Simulation in Electronics
- Simulation in Engineering
- Automotive Simulation
- Simulation in Logistics
- Simulation in Manufacturing
- Chemical Engineering and Energy Simulation
- Simulation in Robotics and Multibody Applications
- Simulation in Aerospace and Traffic
- Modelling and Simulation in the Textile Industry.

More than 80 scientific papers were presented at ISC'2003, which were published in the proceedings of the conference. Besides these contributions, the conference also included a keynote

speech about knowledge discovery in databases and modelling, and two tutorials concerning parallelism in discrete event simulation and modelling and simulation of complex cell spaces. Together with the ISC'2003 conference a workshop on modelling and simulation in the textile industry took place.

The second edition of ISC (ISC'2004) will be held in Malaga, Spain.

Links:

ISC 2003: <http://www.eurosis.org>
ISC 2004: <http://www.eurosis.org>
Communications Dept. at UPV: <http://www.dcom.upv.es>


Please contact:

Ana Pajares, Universidad Politécnica de Valencia
E-mail: apajares@dcom.upv.es

Philippe Geril, EUROISIS - SIM-CENTER, Ghent University, Belgium
E-mail: Philippe.Geril@ugent.be

CALL FOR PARTICIPATION

CLEF 2004



The Cross-Language Evaluation Forum aims at promoting research into the design of user-friendly, multilingual, multimodal retrieval systems by providing an infrastructure for system testing and tuning, comparison of different approaches, and creation of reusable evaluation test-suits.

The CLEF 2004 campaign offers the following system evaluation tracks: Mono-, Bi- and Multilingual Information Retrieval on News Collections. The document collection for CLEF2004 will contain English, Finnish, French, and Russian documents. A common set of topics (ie structured statements of information needs from which queries are extracted) will be prepared in Dutch, English, Finnish, French, German, Italian, Spanish, Swedish, Russian, Japanese and Chinese. Topics in other languages can be supplied on demand.

Mono- and Cross-Language Information Retrieval on Structured Scientific Data

The rationale for this track is to study retrieval in a domain-specific context using the GIRT-4 German/English social science database with controlled vocabularies in English, German and Russian. Topics will be available in English, French, German and Russian.

Interactive Cross-Language Information Retrieval - iCLEF

This year, the interactive CLEF track will study the problem of Cross-Language Question Answering from a user-inclusive perspective.

Multiple Language Question Answering (QA@CLEF)

Mono- and Cross-Language QA systems will be tested. Languages involved are Dutch, French, German, Italian, Spanish, English.

Cross-Language Image Retrieval (ImageCLEF)

This track evaluates retrieval of images described by text captions based on queries in a different language; both text

and image matching techniques are potentially exploitable. Two test collections are available: St Andrews University historical photographic collection; University Hospitals Geneva medical images.

Cross-Language Spoken Document Retrieval (CL-SDR)

This track aims at the evaluation of CLIR systems on noisy automatic transcripts of spoken documents and will address the following problems: bilingual-SDR from Dutch, French, German, Italian, and Spanish; retrieval with/without known story boundaries; use of multiple automatic transcriptions

Participation

Registration is via the CLEF website.

Important Dates

- Registration opens: 15 January 2004
- Data release: from February 2004
- Submission of runs by participants 15 May 2004
- Release of individual results from 15 July 2004
- Workshop: 16-17 September 2004

More information:

<http://www.clef-campaign.org/>

CALL FOR PARTICIPATION

8th ERCIM Workshop

"User Interfaces For All"

Vienna, Austria, 28-29 June 2004

"User-Centred Interaction Paradigms for Universal Access in the Information Society" is the special theme of the 8th ERCIM Workshop "User Interfaces for All" which builds upon the results of the seven previous workshops.

The vision of User Interfaces for All advocates the proactive realisation of the 'design for all' principle in the field of Human-Computer Interaction (HCI), and involves the development of user interfaces to interactive applications and telematic services, which provide universal access and usability to potentially all users.

Areas of interest for which papers are solicited, include, but are not limited to, the following topics:

Inclusive design and evaluation:

- Design methodologies for universal access
- User interface adaptation
- Intelligent context-sensitive processing
- Modelling task execution contexts
- Universal access in software quality models
- Learnability, tutorials, online help
- Evaluation techniques.

Universal access & novel interaction technologies & tools:

- Multimedia / multimodal interfaces
- Novel interaction techniques, devices and metaphors
- Universal access in new interaction paradigms (wearable and ubiquitous computing, tangible interfaces, Virtual Reality)
- Interface architectures, development tools, interoperability
- Universal access in ubiquitous environments
- Universal access scenarios for AmI environments.

Internet and WWW access:

- Accessibility of (public) web sites
- Privacy and security
- Personalized web content delivery
- Universal access in e-services (e-learning, e-government, e-health, etc)
- Online communities for all.

e-Accessibility & Policies:

- National, European, international policies on e-Accessibility
- Standardisation & universal access
- Legislation for universal access
- Economics of universal access.

Important Dates

- 15 February 2004: Deadline for electronic submission of all papers
- 1 April 2004: Conditional notification of acceptance (confirmation will be given upon registration)
- 1 May 2004: Deadline for electronic submission of camera-ready papers
- 31 May 2004: Deadline for registration.

More information:

<http://www.ui4all.gr/workshop2004/>

CALL FOR PAPERS

European Conference on Digital Libraries (ECDL) 2004

University of Bath, UK,
12-17 September 2004

ECDL 2004 is the 4th in the series of European Digital Library Conferences. ECDL has become the major European forum focusing on digital libraries and associated technical, organisational and social issues.

Digital library research brings together a number of disciplines and practitioner communities, providing a stimulating environment for debate and an opportunity for establishing collaboration. ECDL provides a forum for the exchange of ideas between scientific disciplines and fosters joint initiatives with application communities. Involvement of researchers and practitioners from computing and information science disciplines is well established at ECDL. Increasingly these disciplines are engaging in discussion and co-operation with other groups concerned with knowledge management. ECDL 2004 encourages involvement from a wide range of disciplines and looks forward to increasing contributions from usability experts, educationalists, developers of eLearning systems, and by working within the eScience and GRID communities. The continued success of ECDL rests on the exchange of ideas and experience between these groups.

ECDL 2004 welcomes contribution and participation from scholars, researchers, practitioners and policy makers across the range of disciplines related to the development of the digital library in its widest sense.

Topics

Contributions may include but are not limited to:

Strategy, management and policy making:

- Human resources, economics, business models
- Evaluation, metrics, testbeds.

Digital library applications:

- Digital libraries and learning

- Digital libraries and Research Grids/eScience
- Multilingual digital libraries
- Geo-spatial, still and moving images, sound and 3D graphics libraries.

Information architectures and interoperability:

- Infrastructures and middleware
- Semantic Web technologies
- Agent technologies
- Development of Web Services
- Application of open standards & protocols
- Ubiquitous environments and mobile access.

Users of digital libraries:

- Personalisation and agent technologies
- Annotation and recommendation services
- Presentation and visualisation
- Usability, user studies and user-centred design.

Digital repositories:

- Curation and preservation
- Provenance and trust
- Describing and managing collections
- Digital rights management, licences and schema
- Open archives and self archiving.

Knowledge management:

- Knowledge extraction, data and text mining
- Semantic interoperability, ontologies and knowledge representation
- Classification and indexing, thesauri development
- Metadata schema, registries and mapping

Delivery channels:

- Changing models of scholarly communications
- Collaboratories
- Virtual organisations and digital libraries.

Important Dates

- 5 April 2004: Papers, panels, tutorials and workshops submissions due.
- 19 May 2004: Poster and demonstration submissions due.
- 11 June 2004: Final submission date.

More information:

<http://www.ecdl2004.org>

CALL FOR PARTICIPATION

ECAI'04 —16th European Conference on Artificial Intelligence

Valencia, Spain, 22-27 August 2004

The European Coordinating Committee for Artificial Intelligence (ECCAI) and the 'Grupo de Tecnología Informática' of the 'Departamento de Sistemas Informáticos y Computación' of the 'Universidad Politécnica de Valencia' cordially invite you to participate in the 16th European Conference on Artificial Intelligence (ECAI 2004). ECAI provides a public forum for researchers from academy and industry and public organizations, offering them an invaluable occasion to meet and exchange ideas, thus contributing to making real the upcoming 'Information Society'.

The conference includes, from 22-24 August, two days and a half of workshops and tutorials on most advanced topics, as well as the Second STarting AI Researchers Symposium (STAIRS), followed by the three-day technical program of ECAI. Following an established tradition, the Sub-Conference on Prestigious Applications of Intelligent Systems (PAIS) will run in parallel with ECAI. The technical programme will include refereed paper presentations, invited talks by prestigious speakers and poster sessions.

Important Dates

- Paper abstract submission: 13 February 2004
- Full paper submission: 18 February 2004
- Notification of acceptance : 2 May 2004
- Final paper due : 31 May 2004

More information:

<http://www.dsic.upv.es/ecai2004>

CALL FOR PAPERS

The 13th Conference on Mathematics for Industry

Eindhoven, The Netherlands,
21-25 June 2004

The conference will be held at the Eindhoven University of Technology following up ECMI - 2002 in Jurmala, Latvia. In order to have a broader scope of application of mathematics in industry, ECMI 2004 is jointly organized by ECMI and ENBIS (European Network for Business and Industrial Statistics). The conference will be devoted to mathematical and statistical modelling, analysis and simulation of problems arising in a practical context. In particular the following application areas have been chosen:

- Aerospace
- Electronic industry
- Chemical technology
- Life sciences
- Materials
- Geophysics
- Financial mathematics
- Water flow.

Important Dates

- Submission of proposals for minisymposia: 30 January 2004
- Early-bird registration: 1 February 2004
- Submission of papers and posters 25 February 2004
- Notification of acceptance of papers and posters: 25 March 2004.

More information:

<http://www.ecmi2004.tue.nl/>

CALL FOR PARTICIPATION

STAIRS 2004 — Second European Starting AI Researcher Symposium

Valencia, Spain,
23 - 24 August 2004

STAIRS'2004 is the second European Starting AI Researcher Symposium, an international meeting aimed at fostering interaction and exchange between AI researchers, from all countries, at the beginning of their career. STAIRS'04

offers PhD students and people holding a PhD for less than one year:

- a first experience on submitting and presenting a paper in an international forum with a broad scope and thorough selection process
- the opportunity to gather knowledge and exchange ideas related to their research problems and approaches
- information on European research careers and mobility.

STAIRS'04 will be held jointly with ECAI'04 in Valencia, Spain, on August 23rd and 24th. The Symposium will be hosted by the 'Departamento de Sistemas Informáticos y Computación (DSIC) of the Universidad Politécnica de Valencia'.

Important Dates

- 1 April 2004: Title and abstract
- 7 April 2004: Paper submission deadline
- 21 May 2004: Notification of acceptance sent to authors
- 4 June 2004: Camera-ready version due
- 23-24 August 2004: Symposium venue

More information:

<http://www.dsic.upv.es/ecai2004/stairs2004>

CALL FOR PARTICIPATION

Symposium on Point-Based Graphics

Zurich, Switzerland, 2-4 June, 2004

The Symposium on Point-Based Graphics to be held at the Swiss Federal Institute of Technology in Zurich, reflects the current renaissance of graphical methods based on point-sampled geometry; this is a result of the dramatic increase in computational complexity observed when using polygonal models.

Manuscripts should be submitted by 7 April 2004.

Suggested topics

- Data acquisition
- Surface reconstruction
- Representation of point-sampled geometry
- Sampling, approximation, and interpolation
- Digital geometry processing of point models

- Transmission and compression of point-sampled geometry
- Geometric modelling using point primitives
- Rendering algorithms and rendering architectures for point primitives
- Animation and morphing of point-sampled geometry
- Hybrid representations and algorithms.

More information:

<http://graphics.ethz.ch/pbg/cfp.html>

SPONSORED BY ERCIM

MFCS 2004 — 29th International Symposium on Mathematical Foundations of Computer Science

Prague, Czech Republic,
22-27 August, 2004

The series of MFCS symposia, organized alternately in the Czech Republic, Poland and Slovakia since 1972, has a long and well-established tradition. The MFCS symposia encourage high-quality research in all branches of theoretical computer science.

The topics include algorithms and data structures, automata, grammars and formal languages, complexity theory (structural and computational), computational biology, computational geometry, concurrency theory, cryptography and security, databases and knowledge-based systems, foundations of programming, formal specifications and program development, machine learning, models of computation (parallel and distributed computed, fixed parameter tractability, quantum computing, and other), semantics and logics of programs, and theoretical issues in artificial intelligence.

Important Dates

- Submission deadline: 15 March 2004
- Acceptance/Rejection: 14 May 2004
- Final version due: 6 June 2004
- Early registration: 15 June 2004

More information:

<http://mfcs.mff.cuni.cz/>

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

Copyright Notice

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

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

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

Director: Michel Cosnard

Central Editor:

Peter Kunz
peter.kunz@ercim.org

Local Editors:

AARIT: n.c.
CCLRC: Martin Prime
M.J.Prime@rl.ac.uk
CRCIM: Michal Haindl
haindl@utia.cas.cz
CWI: Annette Kik
Annette.Kik@cw.nl
CNR: Carol Peters
carol.peters@isti.cnr.it
FORTH: George Papatzanis
gpap@ics.forth.gr
FhG: Michael Krapp
michael.krapp@scai.fraunhofer.de
FNR: Patrik Hitzelberger
hitzelbe@crpgl.lu
INRIA: Bernard Hidoine
bernard.hidoine@inria.fr
Irish Universities Consortium:
Ray Walshe
ray@computing.dcu.ie
NTNU: Truls Gjestland
truls.gjestland@ime.ntnu.no
SARIT: Harry Rudin
hrudin@smile.ch
SICS: Kersti Hedman
kersti@sics.se
SpaRCIM: Salvador Lucas
slucas@dsic.upv.es
SRCIM: Gabriela Andrejkova
andrejk@kosice.upjs.sk
SZTAKI: Erzsébet Csuhaaj-Varjú
csuhaaj@sztaki.hu
VTT: Pia-Maria Linden-Linna
pia-maria.linden-linna@vtt.fi

Free subscription


You can subscribe to ERCIM News free of charge by:

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

Book Review

The ITU has released the fifth book in its Internet Reports series entitled "The Birth of Broadband". The book is a fountain of statistics and analyses.


The content is best summarizing by quoting from the Foreword: "The introductory chapter of this report, Broadband dreams, explains what broadband can do for users, society and industry. Chapter two, Technologies for broadband, explains the different broadband technologies and how each can provide broadband access under different economic and network conditions. Chapter three, Supplying broadband, looks at how broadband has been successfully provided in certain economies and how certain policies can help expand the network. Chapter four, Using broadband, discusses the current and emerging applications that are driving broadband take-up along with applications and content models that show the most promise for the future. Chapter five, Regulatory and policy aspects, examines regulatory and policy frameworks in successful broadband markets. Chapter six, Promoting broadband, looks at the broadband experiences of several countries characterized by high penetration rates and extensive networks, including conclusions drawn from ITU country case studies on broadband, and examines why and how broadband should be actively promoted. Chapter seven, Broadband and the information society, looks at broadband as a component of a society built around ubiquitous access to information, including some of the benefits and pitfalls of total connectivity. The Statistical annex contains data and charts covering 206 economies worldwide, with original data on broadband and comparative information measured against a selection of variables."



Surprising to this reviewer is how much greater broadband penetration is in some of the Far Eastern countries, particularly in the Republic of Korea which leads the rest of the world. In Korea, over 20 of 100 inhabitants have broadband and over 40 percent of Internet subscribers access the Internet via broadband. Figures for cost and hours of usage are given.

The book, nearly 200 pages, may be ordered via sales@ITU.int or the ITU Web site: <http://www.itu.int/osg/spu/publications/sales/birthofbroadband/index.html>

Harry Rudin, Consultant and Swiss Local Editor for ERCIM News

 **INRIA — Michel Cosnard** became the new president of INRIA following a decree signed by the French President Jacques Chirac. Michel Cosnard takes over from Bernard Larrouturou, who was recently appointed Chief Executive Officer of the French National Centre for Scientific Research (CNRS).





© INRIA / Photo G. Favier - Agence Vu

Michel Cosnard was born in 1952. He received his engineering degree in computer science from ENSIMAG (Grenoble, France), his Master's degree in applied mathematics from Cornell University (USA) in 1975 and his Doctorat d'Etat in computer science from the Université de Grenoble (France) in 1983. He joined the CNRS in 1979, and in 1987 he was appointed professor of computer science at the 'Ecole Normale Supérieure de Lyon' (France), where he founded the 'Laboratoire d'Informatique du Parallélisme' (LIP). He served as chairman of the LIP until 1997.


In 1997, he became director of the research unit of INRIA in Lorraine, and of LORIA ('Laboratoire Lorrain de Recherche en Informatique et ses Applications'). In June 2001, he was nominated director of the INRIA research unit in Sophia Antipolis (Nice). His research interests are in the design and analysis of parallel algorithms, and the complexity analysis of automata and neural nets. In January 2001, he was appointed director of the French National Research Program on Grid Computing by the Ministry of Research. In 1994, Michel Cosnard was awarded a national prize from the French National Academy of Science. In 1995, he received the IFIP Silver Core and in 2003, the Charles Babbage Award. He has published more than 100 papers related to parallel processing.


Michel Cosnard will represent INRIA on ERCIM's Board of Directors and will continue to act as ERCIM's manager.

 **Fraunhofer-Gesellschaft — A Sino-German Mobile-Communications-Institute**, located at the Fraunhofer Institute for Telecommunications, Heinrich-Hertz-Institut in Berlin was established in December 2003. The institute will cover various aspects of mobile communications and also to corroborate and initiate joint Chinese and German research activities. Primary responsibility is both theoretical and practical research work for evolving third generation mobile communication systems. The work is supported by the Ministry of Science and Technology (MOST), P.R. China and the Federal Ministry of Education and Research (BMBF), Germany. The institute is the second wing of the Sino-German research establishment after the 'Sino-German Joint Software Institute' which was opened in Beijing earlier in 2003. sion takes place via a network of antennas at fixed locations, temporary ad-hoc networks can be set up using appropriate technologies. Such temporary networks can be used, for example, for car-to-car communication in vehicle fleets.

 **INRIA — Texas Instruments has established a Java™ competence center in Rennes, France** in December 2003 as an extension of its effort to expand the availability of innovative wireless multimedia applications on TI OMAP™ processors and wireless communications chipsets. The center will open following a five-year research program conducted jointly with INRIA, that was focused on developing technology that will optimize the performance of Java applications when running on cell phones and other portable and converged multimedia electronics. One of the primary goals of the Java competence center will be to drive the transition of the research results into business reality, thus enabling the large community of Java developers an accelerated application performance when they port their multimedia applications onto mobile terminals. During the initial research stage of the relationship, TI and INRIA focused on technology development that addresses the complexity of performance and power-hungry Java applications and optimizes their compat-

ibility with the real-time constraints imposed by several multimedia applications running on the same mobile device. By combining INRIA's expertise in embedded systems software development with TI's wireless systems know-how, the two partners were able to solve the technical challenges of the research stage during which over 30 patents were being generated. The core team of the new TI competence center is comprised of the former INRIA engineers and researchers who worked with TI on the original joint program. The center will act as a key interface for TI's customers and partners in the Java application area and is located in Rennes to benefit from the close proximity to INRIA-Rennes for further joint research activities.

 **NTNU — The Norwegian University of Science and Technology will spend 40 million NOK (5 million Euro) to enhance the research and studies in language technology.** During the last 6-7 years, a number of different projects in language technology have been launched at NTNU. Most of the projects are externally financed, mainly by the Research council of Norway, and their priority programme KUNSTI (development of knowledge in Norwegian language technology). Nevertheless, many of the minor projects are financed by internal funds at NTNU. The 40 millions will be shared between the various projects at NTNU. The main projects are BRAGE (user interface with natural language), which focuses on human-machine dialogue, and the allied project VOCALS, which focuses on the convergence between communication systems, advanced dialogue management and spoken language technology. There is also a project between NTNU and the University in Bergen called LOGON, which is a project devoted to machine translation from Norwegian to English. In addition, there are also many on-going minor projects. An application developed at NTNU is already commercialized and available to the public. This application has a natural language interface, and the system can give you information about the bus schedules in Trondheim. You can either use the web to send your question to the server, or you can choose to use SMS.



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/>



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/>



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/>



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/>



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/>



Spanish Research Consortium for Informatics and Mathematics
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



Czech Research Consortium for Informatics and Mathematics
FI MU, Botanická 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>



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/>



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/>



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



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



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



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/>



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/>



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/>



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



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>



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

Order Form

Name:

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.

Organisation/Company:

Address:

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

Post Code:

City:

Country:

E-mail:

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