


Number 74, July 2008

ERCIM



NEWS

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

Special:

Supercomputing at Work

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of 10,500 copies. The printed version of ERCIM News has a production cost of €8 per copy. Subscription is currently available free of charge.

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

Editorial Board:

Central editor:

Peter Kunz, ERCIM office (peter.kunz@ercim.org)

Local Editors:

Austria: Erwin Schoitsch, (erwin.schoitsch@arcs.ac.at)

Belgium: Benoît Michel (benoit.michel@uclouvain.be)

Denmark: Jens Bennedsen (jbb@it-vest.dk)

Czech Republic: Michal Haindl (haindl@utia.cas.cz)

France: Bernard Hidoine (bernard.hidoine@inria.fr)

Germany: Michael Krapp (michael.krapp@scai.fraunhofer.de)

Greece: Eleni Orphanoudakis (eleni@ics.forth.gr)

Hungary: Erzsébet Csuhaj-Varjú (csuhaj@sztaki.hu)

Ireland: Ray Walsh (ray@computing.dcu.ie)

Italy: Carol Peters (carol.peters@isti.cnr.it)

Luxembourg: Patrik Hitzelberger (hitzelbe@lippmann.lu)

Norway: Truls Gjestland (truls.gjestland@ime.ntnu.no)

Poland: Hung Son Nguyen (son@mimuw.edu.pl)

Spain: Salvador Lucas (slucas@dsic.upv.es)

Sweden: Kersti Hedman (kersti@sics.se)

Switzerland: Harry Rudin (hrudin@smile.ch)

The Netherlands: Martine Roeleveld (Martine.Roeleveld@cw.nl)

United Kingdom: Martin Prime (M.J.Prime@rl.ac.uk)

W3C: Marie-Claire Forgue (mcf@w3.org)

Contributions

Contributions must be submitted to the local editor of your country

Copyright Notice

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

Advertising

For current advertising rates and conditions, see

<http://ercim-news.ercim.org/> or contact office@ercim.org

ERCIM News online edition

The online edition is published at <http://ercim-news.ercim.org/>

Subscription

Subscribe to ERCIM News by: contacting the ERCIM office (see address above) or by filling out the form at the ERCIM website:

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

Next issue:

October 2008, Special theme: Safety-Critical Software

Cover illustration:

Massively parallel simulations of aircraft wake instabilities

An aircraft wake consists of powerful long-lasting trailing vortices. This potential hazard imposes safety distances, and thus limits airport traffic. Thousands-of-processors high-resolution simulations can accurately capture fast growing instabilities which perturb the vortices and can therefore accelerate their decay. The cover shows the volume rendering of vorticity in the case of a fast-growing instability. Secondary vortices generated by the stabilizer reconnect with the wing ones and result in a disturbance that propagates along and inside the vortex cores.

Credits:

Philippe Chatelain, Michael Bergdorf, Diego Rossinelli,

Petros Koumoutsakos, ETH Zurich, Switzerland

Alessandro Curioni, Wanda Andreoni,

IBM Zurich Research Laboratory, Switzerland.

Acknowledgments:

IBM T.J. Watson Research Center

Swiss Supercomputing Center.

Knocking at Petaflop's Door

Just eleven years ago, the first supercomputer, Intel's ASCI Red, broke the barrier of more than one trillion (short-scale number system) floating point operations per second: one Teraflop/s. Last month, a new system at the Los Alamos National Lab, nicknamed 'Roadrunner', achieved for the first time the staggering result of one Petaflop/s or one quadrillion floating point operations per second. Roadrunner was benchmarked through the Linpack benchmark code, which is used to compare and rank the performance of the 500 fastest supercomputers worldwide (<http://www.top500.org>). The enormous 1000-fold increase in performance of supercomputers in the last eleven years highlights the dramatic progress of these universal instruments. Supercomputers are the base on which has been built the continuous advancement of computer simulations, now a key element in the scientific and industrial competitiveness of knowledge-based economies in the 21st century. Simulations are the engines for industrial fields such as aeronautics, and the automotive, pharmaceutical, oil and financial industries. They drive progress in key scientific fields of the highest societal relevance, like climatology, fusion energy and biology, not to mention the defence sector with its requirement for reliability testing and maintenance of nuclear weapons without the use of nuclear testing.

It is not by chance that it is an American system that has made the enormous leap over the Petaflop/s mark: the USA recognized the growing relevance of simulation science at an early stage. In 1991, they passed the High-Performance Computing Act, which states that HPC be given top priority for research. This initiative had a very strong impact on the coordination of all programs with dedicated budgets for supercomputing, and secured the US's leadership in the field. Since then, more than half of the fastest systems worldwide have been located in the US. The importance of HPC has also penetrated all levels of the federal administration, and various presidents have declared their strong engagement in this sector on a number of occasions.

Parallel to the USA, which started the race to Petaflop/s around 1997, Japan has also pursued a very active policy in support of HPC. Developed and integrated by NEC in Japan, the 'Earth Simulator' was the most powerful supercomputer in the world from 2002 to 2004. Japan's next-generation supercomputer project aims at delivering 10 Petaflop/s in the year 2012. In addition, China recently announced that it will design and build new supercomputers and join the leading countries in 2010 with the installation of a top-level system.

Europe's decision makers have so far placed supercomputing for simulations in science and engineering on a much lower level of priority. Europe's previous framework programme, FP6, concentrated its efforts mostly on embedded systems, telecommunications, distributed computing and data services.

Fortunately there is good reason to be optimistic that the simulation sciences – often denoted as the third column of knowl-



*Prof. Dr. Dr. Thomas Lippert
Director of the Institute for Advanced Simulation,
Head of Jülich Supercomputing Centre.*

edge creation – will also flourish over here. Support for supercomputing infrastructures has been established in the 7th Framework Programme, and following the late 2007 recommendations of the European Strategy Forum on Research Infrastructures, the European Commission has put the creation of a European supercomputing infrastructure on its agenda. These top-level systems will provide levels of performance comparable to other top installations worldwide for the benefit of Europe's scientific and technological development.

In spring 2007, fourteen European countries formed the initiative 'Partnership for Advanced Computing in Europe' (PRACE), in order to be able to respond to the Commission's call for proposals. The PRACE project was granted by the European Commission in the autumn of 2007 and commenced in January 2008 (see <http://www.prace-project.eu>). It is coordinated by the German Forschungszentrum Jülich. Its primary objective is the creation of an organizational and legal framework for a European supercomputing infrastructure by 2010. The new 'tier-0' systems (five European countries – the principal partners France, Germany, Spain, the Netherlands and the United Kingdom – announced their support for the infrastructure through the installation of tier-0 systems) will establish a new level of performance in Europe, built of the national HPC infrastructures and the European Grids.

The second important mission of PRACE is to create in Europe a framework for the development of the next generation of supercomputers and revitalize the European supercomputer industry. Leading HPC companies in Europe are keen to contribute, and together with research institutes and universities have formed the consortium PROSPECT in order to cooperate with PRACE. Their vision is to create a European technology platform for supercomputing.

The first signs are now visible of a resurgence of European systems in the top500 list. With the help of PRACE, Europe will soon be knocking at Petaflop's door, and Europe's industry may in fact become strong enough to pave the way towards Exaflop performance.

Thomas Lippert

2 Editorial Information

KEYNOTE

- 3 **Keynote**
by Thomas Lippert
Director of Institute for Advanced Simulation,
Head of Jülich Supercomputing Centre

JOINT ERCIM ACTIONS

- 6 **Euro-India ICT Cooperation Initiative**
by Florence Pesce
- 7 **ERCIM "Alain Bensoussan" Fellowship Programme**
- 8 **D4Science - Deploying Virtual Research Environments**
by Donatella Castelli and Jessica Michel
- 9 **Advisory Committee for ERCIM**
- 10 **First CRM-INRIA-MITACS Workshop**
by Mark Thiriet

THE EUROPEAN SCENE

- 12 **New Research Perspectives in Engineering Secure Complex Software Systems and Services**
by Thomas Skordas
- 13 **EC-ERCIM Strategic Seminar on ICT Security: "Engineering Secure Complex Software Systems and Services", Brussels 16 October 2008**

SPECIAL THEME

Introduction to the Special Theme

- 14 **Supercomputing at Work**
by Alessandro Curioni and Ray Walshe
- Weather and Environment**
- 16 **Quantitative Seismic Monitoring Methods**
by Martin Landrø
- 17 **Studying CO₂ Sequestration with the Power of Supercomputing**
by Klaus Johannsen, Andreas Kopp, Olli Tourunen and Josva Kleist
- 19 **Numerical Weather Forecasting for Poland**
by Maciej Szpindler and Maciej Cytowski
- 20 **Supercomputing: Weathering a Changing Environment**
by Ray Mc Grath
- Technology**
- 22 **Supercomputing at Work in the nanoCMOS Electronics Domain**
by Richard Sinnott, Campbell Millar and Asen Asenov
- 23 **SPRANKER: A Discovery Tool to Rank Service Providers Using Quality of Experience**
by Domenico Laforenza, Franco Maria Nardini, Fabrizio Silvestri and Gabriele Tolomei
- 25 **Design of a Functional Architecture for the Management of Cluster Resources and Services through the Web**
by Juan Antonio Ortega, Jorge Cantón, Ana Silva, David Bosque and Francisco Velasco
- 26 **Exploitation of Cell Multi-Processor Array in Solution of Spatio-Temporal Dynamics**
by Zoltán Nagy, László Kék, Zoltán Kincses, András Kiss and Péter Szolgay
- 27 **IANOS – Efficient Use of HPC Grid Resources**
by Philipp Wieder, Wolfgang Ziegler and Vincent Keller
- Medicine**
- 29 **Towards Personalized Medicine: High-Performance Computing in the Life Sciences**
by Olaf Schenk, Helmar Burkhart and Hema Reddy

31 Microstructural Finite Element Analysis of Human Bone Structures

by Peter Arbenz and Ralph Müller

32 Using Desktop Grids to Securely Store e-Health Data

by Jesus Luna, Manolis Marazakis and Marios D. Dikaiakos

34 Supercomputing in Clinical Practice

by Stefan Zasada, C.V. Gale, Steven Manos and Peter Coveney

[Large-Scale Simulation](#)

35 Large-Scale Immune Models and Visualization

by Dimitri Perrin and John Burns

37 High-Performance Computing for Modelling Bacterial Communities

by James T. Murphy, Ray Walshe and Marc Devocelle

38 Harnessing the Power of Supercomputing

by Richard Blake

40 Fluids and Supercomputers: The Billion Particle Era

by Petros Koumoutsakos

41 Supersonic Flow Simulation on Emulated Digital Cellular Neural Networks

by Sándor Kocsárdi, Zoltán Nagy, Árpád Csík and Péter Szolgay

43 High-Performance Computing of Multiphase Flow

by Bipin Kumar, Yan Delaure and Martin Crane

44 Set-Top Supercomputing: Scalable Software for Scientific Simulations on Game Consoles

by Dimitrios S. Nikolopoulos

R&D AND TECHNOLOGY TRANSFER

46 Enhancing Traffic Safety by Integrating Real-Time Infrastructure and Vehicle Data in a Cooperative System

by Kashif Din

47 Realizing Ambient Assisted Living Spaces with the PERSONA Platform

by Francesco Furfari and Mohammad-Reza Tazari

49 Medical Record Keeping Made Visually Accessible

by Harry Rudin

50 Games on Networks

by Jacek Miękisz

51 Carrot²: Making Sense of the Haystack

by Stanisław Osiński and Dawid Weiss

52 Quantifying WiMAX Performance

by Kostas Pentikousis, Ilkka Harjula, Esa Piri and Jarno Pinola

54 DEPLOY: Industrial Deployment of Advanced System Engineering Methods for High Productivity and Dependability

by Alexander Romanovsky

EVENTS

56 LREC 2008 - The Language Resources and Evaluation Conference

by Nicoletta Calzolari and Khalid Choukri

57 DELOS Summer School on Preservation

by Vittore Casarosa

58 The Fuschl Conversations – A Central Activity of Systems Sciences

by Gerhard Chroust

58 Announcements

62 In Brief

Next issue: October 2008

Next special theme: Safety-Critical Software

Euro-India ICT Cooperation Initiative

by Florence Pesce

ERCIM is a partner in the Euro-India ICT Cooperation Initiative ('EUROINDIA'), a Support Action under the Seventh Framework Programme. It aims to reinforce collaborative research and innovation activities between Indians and Europeans in various ICT domains targeted under FP7.

EUROINDIA benefits from earlier groundbreaking work done under BASIC (EuroIndia2004 Cooperation Forum on the Information Society) and MONSOON (<http://www.euroindia-it.org>): projects which identified, profiled and engaged the widespread Indian ICT research community and helped to significantly increase Indian participation in FP6 and FP7 projects.

EUROINDIA aims to move further down this promising road and sustain collaborative EU-India research and technology development (RTD) potential in a durable manner.

The key objectives of the project are:

- to extensively map ICT research and innovation activities across India
- to perform an in-depth survey of Indian ICT R&D players
- to support the policy dialogue between the EU and India in ICT domains.

The mapping and survey activities will reveal the state of ICT research and innovation in India while identifying suitable themes and personnel for collaborative projects with EU research establishments under FP7 and beyond.

The annual policy dialogue between the European Commission and India will be enriched and strengthened by including the views of ICT research communities from both regions. This will also help in the translation of policy recommendations and joint action agendas into concrete cooperative projects.

During its two-year duration (January 2008 to December 2009), the project will organize six Technology Brainstorming workshops, six ICT information days and two international conferences in India. These events will foster networking within the ICT communities, leading to the identification of mutual areas of interest and development of joint EU-India research projects.

The EUROINDIA kick-off meeting, hosted by the Copenhagen Business School (CBS), was held in Copenhagen on 20-21 January 2008. The meeting started with an opening conference in the presence of the CBS President, Mr Finn Junge-Jensen and the Indian ambassador to Denmark, Sri (Mr) R S Ray. Professor Mogens Kühn Pedersen, head of the department of Informatics at CBS and EUROINDIA project coordinator, introduced the project and presented a brief overview of the scope and objectives of the EUROINDIA



Participants of the kick-off meeting, held in Copenhagen on 20-21 January 2008.

initiative. Professor Krithi Ramamritham, Dean of R&D at the Indian Institute of Technology, Bombay (IIT-B) and Dr Gérard Huet, Senior Research Scientist at INRIA (France), spoke on ICT R&D practices and approaches from Indian and European perspectives. The Indian ambassador focused on the scope and potential of research and development between India and Denmark.

As planned during the kick-off meeting, the first EuroIndia Information Day was held on 22 April 2008 and hosted by the Indian Institute of Sciences in Bangalore. It offered participants an opportunity to learn how to collaborate with European entities on competitive ICT RTD projects funded by the European Union, by providing them with concrete examples and advice.

The first part of the day was dedicated to a presentation by the European Commission Delegation to India on the state of EU-India science and technology cooperation, with particular emphasis on actions in the ICT domains. EUROINDIA project partner Infra Technologies followed up with a presentation that highlighted the various aspects, instruments, rules and realities pertaining to Indian participation in FP7 projects.



Professor Jamadagni, Chairman, Centre for Electronic Design & Technology (CEDT), Indian Institute of Sciences (IISc), Bangalore at the first EuroIndia Information Day held in Bangalore on 22 April 2008.

Successive sessions conducted by the Copenhagen Business School on internal innovation in organizations in India and on collaborative ICT R&D saw lively interaction between participants and the panel, bringing to the fore issues having a critical bearing on collaborative research projects. Another project partner, the Indian Institute of Technology – Bombay (IIT-B), gave its own perspective on EU-India collaborative research in ICT and shared its experience of participation in EU-funded projects.

Fifty-five participants attended the first information day, 13% being from Europe and 87% from India. Of these attendees, 47% came from academia and research, 20% from small and medium enterprises (SMEs), 16% from consultancies and the remaining 17% from government/public administration, non-profit organizations and research labs.

The EUROINDIA project is coordinated by the Copenhagen Business School – Department of Informatics (CBS-INF), with the following partners from the EU and India: Infra Technologies Sarl (INFRA), France; Trust-IT Services Ltd, United Kingdom; the Federation of Indian Chambers of Commerce & Industry (FICCI) India; the Indian Institute of Technology – Bombay (IIT-B), India; and ERCIM.

Link:

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

Please contact:

Florence Pesce

ERCIM Office

Tel: +33 4 92 38 50 73

E-mail: florence.pesce@ercim.org

ERCIM "Alain Bensoussan" Fellowship Programme

Who can apply?

ERCIM offers fellowships for PhD holders from all over the world. ERCIM encourages not only researchers from academic institutions to apply, but also scientists working in industry.

Why to apply?

The Fellowship Programme enables bright young scientists from all over the world to work on a challenging problem as Fellows of leading European research centres. The programme offers the opportunity:

- to work with internationally recognized experts
- to improve their knowledge about European research structures and networks
- to become familiarized with working conditions in leading European research centres
- to promote cooperation between research groups working in similar areas in different laboratories.

What is the duration?

Fellowships are generally of 18 month duration, spent in two of the ERCIM institutes. In particular cases a Fellowship of 12 month duration spent in one institute might be offered.

Which topics?

The Fellowship Programme focuses on topics defined by the ERCIM Working Groups and projects managed by ERCIM such as biomedical informatics, computing and statistics, constraints technology and applications, embedded systems, digital libraries, environmental modelling, e-mobility, formal

Photo by courtesy of INRIA



methods, Grids, security and trust management, software evolution, and many other areas in which ERCIM institutes are active.

Where are the Fellows hosted?

The Fellows are hosted by ERCIM member institutes. Nineteen leading European research organisations are currently members of ERCIM.

What are the conditions?

- Have obtained a PhD degree during the last 4 years (prior to the application deadline) or be in the last year of the thesis work with an outstanding academic record
- be fluent in English
- be discharged from military service
- start the grant before October 2009
- have completed the PhD before starting the grant (a proof will be requested).

In order to encourage mobility, a member institute cannot host a candidate of the same nationality. A candidate cannot be hosted by

a member institute, if by the start of the Fellowship, he or she has already worked in this institute for a total of six months or more, during the last three years.

Competitive salary

For the entire period, the Fellow is paid an attractive monthly allowance which may vary depending on the country.

How to apply ?

Applications have to be submitted online. The application form will be available one month prior to the application deadline at the URL given below.


Next deadline for applications
30 September 2008

More information
<http://www.ercim.org/activity/fellows/>

D4Science - Deploying Virtual Research Environments

by Donatella Castelli and Jessica Michel

The D4Science (Distributed colLaboratories Infrastructure on Grid-Enabled Technology for Science) project aims to continue the path that the GÉANT (a multi-gigabit pan-European data communications network), EGEE (Enabling Grids for E-science in Europe), and DILIGENT (A Digital Library Infrastructure on Grid Enabled Technology) projects have initiated towards establishing networking, Grid-based, and data-centric e-Infrastructures. These e-Infrastructures are expected to accelerate multidisciplinary research by overcoming several crucial barriers that stand in the way, primarily those related to heterogeneity, sustainability and scalability. In order to achieve this objective a D4Science production quality e-Infrastructure has been created, and will be progressively enriched and consolidated. It will provide facilities for creating Virtual Research Environments based on shared computational, data and service resources offered by many different providers like EGEE and large international organizations.



A large aquamarine-coloured phytoplankton bloom is shown stretching across the length of Ireland in the North Atlantic Ocean in this image, captured on 6 June 2006 by Envisat's Medium Resolution Imaging Spectrometer (MERIS), a dedicated ocean colour sensor able to identify plankton concentrations. Harmful algal blooms (HABs) is particularly dangerous to fish farms because the fish cannot flee affected areas. D4Science can, for example, provide the infrastructure for an early warning system of HABs to prevent fish farmers from losing an entire stock in a single day.

Photo: European Space Agency.

The D4Science project builds on the experience of a predecessor testbed project, DILIGENT (A Digital Library Infrastructure on Grid Enabled Technology). This was the first project to propose the integration of digital library and grid technology in order to create an innovative type of e-infrastructure capable of supporting scientific cooperation based on the managed sharing of a variety of resources. The core notions of the solution proposed by DILIGENT are virtual organizations (modelling sets of users and resources and clearly defining what is shared, who is allowed to share, and the conditions under which sharing occurs), and virtual research environments (representing frameworks of applications, service and data sources dynamically identified to support cooperation processes). By implementing these two notions, the DILIGENT e-infrastructure provides a framework for improving synergies across currently fragmented scientific communities, thus aiding research on a global scale.

User Communities

D4Science aims to consolidate and exploit the DILIGENT e-infrastructure to build VREs that will serve the needs of broad scientific disciplines. In its first stage, the focus is on serving the domains of environmental monitoring, and fisheries and aquaculture resource management. The project's user communities – the European Space Agency (ESA), departments of the Food and Agriculture Organization of the United Nations (FAO), and the World Fish Center, a member of the Consultative Group on International Agricultural Research – will benefit greatly from access to shared resources. The e-infrastructure will, for example, provide them with a powerful instrument for ecosystem modelling. This application requires satellite data (eg ocean colour and reef maps), climate data, hydrographic data and a large variety of other types of data (eg environmental reports, species taxonomies, time series etc), as well as specific, often computationally intensive tools for processing this data.

Cooperation between these partners will also encourage the use of data streams from diverse scientific communities to perform socio-ecosystem modelling. The work that the two communities wish to undertake can be applied to the whole domain of biodiversity management (natural resources management), biodiversity conservation, and biodiversity exploitation as catch (mainly fisheries) and farming (both land and aquatic, ie aquaculture). The project will therefore bring together not only the typically separated communities of fisheries/aquaculture and environmental monitoring, but will also create a framework and a process to promote further cooperation of this type across other related communities. Technological Challenges D4Science addresses many important technological challenges. In particular, it will:

- Consolidate and enhance the technology which underpins the D4Science e-infrastructure operation, namely the gCube framework. gCube, successfully deployed within the testbed developed by the DILIGENT project, reflects within its name a three-sided interpretation of the grid vision of resource sharing: sharing of computational resources, structured data and application services. As such, gCube embodies the defining characteristics of computational grids, data grids and virtual data grids. More precisely, it builds on gLite middleware (developed by the

EGEE project) for managing distributed computations and unstructured data, includes dedicated services for managing data and metadata, and offers a novel approach for managing application services. Rather than interfacing with the infrastructure, the software that implements the services is literally handed over to it, so as to be transparently deployed across its constituent nodes according to functional constraints and quality-of-service requirements. This is genuinely ambitious and entirely novel: like computational resources and data before, application logic in gCube becomes a pervasive commodity within an infrastructure that abstracts over the physical location of its resources at any point in time. D4Science now consolidates and enhances the gCube services to reflect the shift in functional and QoS requirements, which marks the passage from a testbed to a production-level infrastructure.

- Make available and maintain a stable, reliable and fully operational e-infrastructure that offers secure sharing and dynamic allocation of resources, where the notion of resources is extended to content and services in addition to computing and storage. By progressively integrating new sites and collecting feedback and requirements, the procedures and policies for managing and operating this infrastructure will also be clarified and consolidated. D4Science will thus become one of the first fully operational infrastructures for virtual data grids.
- Empower the mechanisms for the definition of virtual research environments (by introducing simplified procedures for resource discovery and selection) and their management (by implementing transparent bootstrapping, dynamic configuration and recovery mechanisms for stateful resources; more sophisticated load-balancing mechanisms during resource deployment and discovery; more flexible notification brokering mechanisms and

update policies for resource change; multi-view monitoring mechanisms for resource administration; and finer-grained authentication and authorization policies for resource sharing).

D4Science is one of the main European Research e-Infrastructure projects, supported by the European Commission's Seventh Framework Programme for Research and Technological Development. ERCIM manages the project's administrative and financial activities, while ISTI-CNR leads the scientific coordination. The project started in January 2008 and will have a duration of two years.

It is expected that the D4Science e infrastructure will have a multiplicative benefit to many scientific fields, and will also act as a catalyst for the kind of cooperation and cross-fertilization among multiple communities that is necessary to address many of the grand challenges of science and society.

Links:

- D4Science: <http://www.d4science.eu/>
- gCube: <http://www.gcube-system.org/>
- gLite Lightweight Middleware for Grid Computing: <http://glite.web.cern.ch/glite/>
- EGEE: Enabling Grids for E-science in Europe: <http://public.eu-egge.org>

Please contact:

Donatella Castelli, project director
ISTI-CNR, Italy
E-mail: donatella.castelli@isti.cnr.it

Jessica Michel, administrative and financial director
ERCIM office
E-mail: jessica.michel@ercim.org

Advisory Committee for ERCIM

The ERCIM Board of Directors (BoD) has invited a panel of international experts form Advisory Committee for ERCIM. The committee meets once a year and reports to the ERCIM president and hence to the Board of Directors. The committee was set up to:

- advise the ERCIM President and BoD on overall strategic directions in ICT and applied mathematics including the portfolio of the ERCIM Working Groups
- evaluate Working Group performance
- evaluate the final stage of selection of the Cor Baayen Award
- comment on any ERCIM publications such as strategic reports for the European Commission as requested by the BoD
- propose changes to any aspect of the operation of ERCIM which is beneficial

- undertake any other tasks requested of it by the BoD.

The Advisory Committee reports to the President and hence to the BoD. The current members of the ERCIM Advisory Committee are:

G rard Berry

Chief Scientist, Esterel Technologies, France, Professor at Coll ge de France

Michael L. Brodie

Computer Scientist, USA

Chris Horn

IONA Technologies, Ireland

Frank Kelly

Professor of the Mathematics of Systems in the University of Cambridge, UK

Kurt Mehlhorn

Director of Max Planck Institute for Informatics, Germany

Erik Sandewall (Rapporteur)

Professor of Computer Science at Link ping University, Sweden

Alexander "Lex" Schrijver

Mathematician at CWI and Professor at the University of Amsterdam, The Netherlands

Ulrich Trottenberg

Director of the Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Germany

Mazin Yousif (Chair)

Chief Technology Officer, Avirtec Corporation, USA.

ERCIM IM2M Working Group Event

First CRM-INRIA-MITACS Workshop

by Mark Thiriet

Members of the ERCIM Working Group on 'IT and Mathematics Applied to Interventional Medicine' (IM2IM) recently participated in the first joint workshop of the Centre de recherches mathématiques at the Université de Montréal (CRM), INRIA and the Canadian Network of Centres of Excellence for the Mathematical Sciences (MITACS). This workshop was held in Montreal, Canada on 5-9 May, 2008.

The workshop focused on six main topics in medical imaging, modelling and simulation: 'Heart', 'Physiological Flows', 'Brain Imaging and Function', 'Composite Media, Multiscale Modelling, Optimal Shape and Design', 'Biological Systems, Cell and Tissues' and 'Implantable Medical Devices, Drug Delivery'. The workshop presentations are briefly summarized below, while a detailed description of the presentations is available at the URL provided at the end of the article.

Francois Fages (INRIA) spoke on formal cell biology using BIOCHAM software, and emphasized (1) rule-based modelling of biochemical reaction systems; (2) temporal logic formalization of biological properties; and (3) search algorithms for kinetic parameter values. Elementary rules include complexation/decomplexation, phosphorylation/dephosphorylation, synthesis, transport and degradation.

Céline Grandmont (INRIA) talked on multiscale modelling of the respiratory tract, focusing on the thoracic airways from trachea to terminal alveolar duct. Her work is aimed at developing simple but representative models of small and mid-sized intraparenchymal airways either of the entire distal tracheobronchial tree or after compartment splitting. The simple flow model is coupled to a 3D Navier-Stokes model of incompressible air flow in large thoracic airways. Yves Bourgault (Ottawa University) carried out simulations of aerosol convection in airways using a compressible fluid solver.

André Garon (École Polytechnique de Montréal) and Marie-Isabelle Farinas (Université du Québec à Chicoutimi) presented the work carried out on a patient-specific model of steady flow in the carotid artery network in the context of implantation of a small (hence continuously operating mode) ventricular assist pump (VAD).

Jacques Bélair (Université de Montréal) gave a modelling example on erythropoiesis with age structure and variable life span. Erythrocyte (RBC) production is about 1012 cells per day, and production in various cellular compartments of erythropoiesis from stem cells and different progenitors has a given duration. RBC lifespan (120 days) is also incorporated into the equation. Conservation rules lead to PDEs that can be solved by the method of characteristics with given formulations of the degradation rate. Jean Clairambault (INRIA) works with Benoit Perthame on mathematical modelling of cell proliferation and its control, focusing on

nanoprocesses within the cell to improve drug efficiency by taking into account the circadian cycle.

Irène Vignon-Clementel's (INRIA) talk focused on fluid-solid interaction in a porous medium, applied to heart perfusion. This work is part of the CardioSense3D INRIA initiative action. Solid grains and the fluid domain (viscous fluid governed by the Darcy law) are coupled. Heart matrix deformation is iteratively coupled to perfusion. Myocardium contraction transfers blood from the arterial to the venous compartment. A model of the large coronary artery network is coupled to the perfusion model, using a master interface between a Darcy and Navier-Stokes solver. Dominique Chapelle (INRIA) presented work from the CardioSense3D project, focusing on modelling of the heart's electromechanical coupling. Maxime Sermesant (INRIA) and co-workers have developed fast software for model-based diagnosis and therapy planning for heart diseases. Cardiac image processing combines various approaches, such as deformable model (LV), statistical shape model (RV), interactive segmentation (epicardium) and deformation tracking.




Pressure field of a steady flow in a left coronary network (rigid wall; motionless vessel axis) using fluid solver LifeV. Pressure inlet condition; traction-free outlet condition to be coupled to a perfusion model in a porous medium (from JF. Gerbeau, G. Rossi and I. Vignon-Clementel).

Raymond Spiteri (University of Saskatchewan) carried out a simulation of electrical heart activity in the framework of the virtual heart project. Ionic motions in excitable cardiomyocytes are represented by ordinary differential equations. The numerical model is based on operator splitting and specific IMEX-Rung-Kutta methods. Najib Zemzemi (INRIA) presented simulations of electrocardiograms. Alexandra Franchitti (INRIA) and Yvon Maday are studying optimization aspects for pacemaker treatment. Investigation of the effects of parameters of the fast-slow dynamics equation shows that reaction parameters affect both reaction and diffusion, whereas diffusion parameters only influence diffusion. Youssef Belhamadia (University of Alberta) proposed a numerical bidomain model of electrical waves in the heart.

Michel Sorine (INRIA) proposed a new signal processing method to assess arterial blood pressure. The analysis uses a scattering-based method (SBSA) to target pressure pulses with solitons.

Matteo Astorino (INRIA) works on fluid-structure interaction down to large contact with application to heart valves. Maurice Doyle (University of Ottawa) has studied blood flow in the presence of assist devices, and is now investigating blood flow coupled to deforming myocardium using Adina software. N. Morcos (INRIA) works on a reduced basis method for blood flow simulations in porous tissues (microcirculation). Two-scale homogenization is not periodic, and simulations require the resolution of a large number of parameterized cells. A reduced basis method based on weak formulation and relying on the Galerkin method allows computations to be sped up.



Streamlines at peak expiratory flow in an image-based model of the tracheobronchial tree (surface mesh derived from the French national 'RMOD' project) coupled to lumped parameter models of small bronchi/bronchioles and pulmonary acini (piston-like motion) using an in-house fluid solver. Resistance is higher in a selected lobule (from L. Baffico, C. Grandmont, M. Grasseau and B. Maury).

In collaboration with Leipzig University, Dirk Drasdo (INRIA) develops models of multicellular tissues, in particular tumour growth and liver regeneration. Tumour modeling must bridge large time and length scales, from a few cells to whole populations. Monolayer cell growth relies on cell-based models. Cell motion is governed by a stochastic equation, the random force associated with cell exploration of its environment being defined by statistical features.

G rard Plante (Universit  de Sherbrooke) spoke on pharmacokinetics with a focus on drug delivery to the interstitial fluid compartment that must be crossed to reach cells. Fahim Nekka (Universit  de Montr al) studies pharmacological intra/intervariability induced by patients' irregular drug intake. The adopted strategy allows pharmacokinetics and dynamics to be easily updated. Annabelle Ballesta (INRIA) and Jean Clairambault are studying the pharmacokinetics and dynamics of anticancer drugs in the cell, and are developing a model for anticancer drugs.

Christian Barillot (INRIA) presented data on neuroinformatics in the context of diseases of the central nervous system to develop computer-aided medicine and surgery tools. The work is being partially carried out by a team from INRIA and McGill University. New surgery rooms are equipped with computer infrastructure to integrate data from sensors, to control effectors, and to feed databases. Image quality is improved by a fast nonlocal method for image restoration and noise reduction that is adapted for various acquisition techniques and modes. Olivier Rousseau (Universit  d'Ottawa) is working on an iterative active contour algorithm applied to heart segmentation. Maureen Clerc (INRIA) processes magnetoencephalography (MEG) signals to explore brain activity. Brain neurons undergo cycles of depolarization/ repolarization.

Jiri Patera (Universit  de Montr al) is working on Fourier-like transforms of digital data given on bound lattices of any dimension, symmetry or density. This technique can be applied to interpolate motions between two frames. Jiri Hrivnak (Universit  de Montr al) is working on two-dimensional (anti)symmetric multivariate exponential functions and corresponding Fourier transforms with Jiri Patera. In the same team, Matthieu Voorons (Universit  de Montr al) uses image interpolation based on Lie group theory and has developed an algorithm for block decomposition of images. Comparisons with classical methods show that there are fewer interpolation errors. The proposed method is faster and more accurate than traditional interpolation algorithms.

Fr d ric Lesage ( cole Polytechnique de Montr al) gave a talk dealing with optical imaging of the spinal cord, and explained how a neuronal activity map can be obtained during visual excitations. Mathieu Dehaes ( cole Polytechnique de Montr al) works with Fr d ric Lesage on stimulation of the brain cortex by visual excitations. Inverted curves are found that have a vascular origin. Jean-Marc Lina ( cole de Technologie Sup rieure) is working on the processing of optical imaging signals.

Robert Owens (Universit  de Montr al) works on blood rheology and has developed a non-homogeneous constitutive model. Andr  Fortin (Universit  Laval) also focused his talk on viscoelastic fluid flow and free surface problems that require unsteady, anisotropic, adaptive remeshing. Michel Fortin (Universit  Laval) presented a numerical procedure used to treat frictional contact in solid mechanics.

Jean Paul Zol sio (INRIA) spoke on hidden regularity by extractor for Neuman problems, after pointing out analogies under certain circumstances between the Maxwell equations, elasticity and wave equations using adequate formulation and boundary conditions in incompressible media. Michel Delfour (Universit  de Montr al) derived new equations for calculating dosage under pulsed flow conditions, which can be used in the design of coated stents to prevent restenosis.

Link:

<http://www-rocq.inria.fr/im2im/>

Please contact:


Marc Thiriet, IM2IM Working Group coordinator
INRIA, France
E-mail: Marc.Thiriet@inria.fr

New Research Perspectives in Engineering Secure Complex Software Systems and Services

by Thomas Skordas

A report from a meeting that was recently held at the European Commission on engineering secure complex software systems and services highlights the present status of this discipline and identifies opportunities for further research work.

A meeting on "Engineering Secure Complex Software Systems and Services" was convened in Brussels on 23 April 2008. The event was part of the preparation activities for the FP7 ICT work programme for 2009-2010 and its subsequent calls for R&D proposals in the area of ICT Security and Trust. It brought together a small group of stakeholders from industry and academia. The objective was to discuss what are the main industrial practices and research advances today and to identify the future challenges and how to address



Report of the consultation meeting on "Engineering Secure Software Systems". The full report and the presentations are available for download at the Cordis Web site (see URL at the end of the article).

them. This article briefly describes the main findings from the meeting and the way forward. The full report and the presentations given by the participants can be found at the Web link provided below.

Today's Industrial Practices and Research Work

In industry, software development is still a human intensive and error prone process, despite the existence of sophisticated development environments and formal method techniques. The success of software delivery is still measured against available budget and time constraints but not against the fulfilment of security requirements. There exist many tools for assessing software security but their results can differ considerably. Progress in verification and validation techniques and in static analysis is tangible, but such techniques are far from being cost-effective and flexible and are not yet widely used by industry. In fact, we still do not know how to measure and assess security, how to benchmark security of components or systems and how to prove security to third parties through verifiable evidence. Furthermore, security needs to follow and be aligned with business models and new business scenarios introduce security issues that cannot be addressed by existing solutions.

On the side of research, there is a multitude of research work underway, as for example in model-based design and model-driven security architectures, proof environments, security verification and validation, risk assessment, etc. At present, secure software engineering approaches cope well with size-limited and well-defined systems having relatively stable requirements. Future IT systems will be at least one order of magnitude more complex and larger than today's systems, evolving more dynamically and necessitating coordination of multiple intervening organizations. This brings new challenges to secure software engineering: How to scale it up to much larger systems with effective cross-organisational governance that aligns with various IT and enterprise governance objectives, while maintaining a sufficiently agile security architecture and configuration to accommodate changing requirements, contexts and use-cases? How to scale it down for the pervasive embedded systems (eg sensors and actuators)?

Shaping the Future

Today, there is a gap between the pragmatic practices of secure software development used by industry and the results of available research. Industry needs agile cycles of software delivery where time-to-market is often a driving force, while academia has often been focusing on time-consuming and heavyweight approaches. Secure software engineering can only be widely adopted by the software industry if there is a clear added value for the business. Progress in the following areas of work can act as a strong incentive for the further consideration and take-up of advanced secure software engineering practices by industry:

- The further development of models, metrics, tools and risk analysis and management processes; risk assessment is a key factor in capturing, specifying, implementing and monitoring security in software systems; the use of common experimental facilities and test data sets for testing and benchmarking new solutions; promoting best practices; and, further addressing the certification of software

– both the source and executable codes – from the security perspective.

- Other areas where further research is needed include: methods and tools for ensuring system security by design, for assessing the security properties and for further automating system assurance (eg, by scaling up verification and validation), the composition of security properties at run time, embedding privacy in software development practices or integrating secure software engineering practices into the corporate IT governance of organisations.

Next Steps

The next event on this topic is the EC-ERCIM Strategic Seminar on ICT Security: "Engineering Secure Complex Software Systems and Services" that will take place in Brussels on 16 October 2008 (see announcement below).

Disclaimer: The content of this paper is the sole responsibility of the author and in no way represents the view of the European Commission or its services.

Link:

http://cordis.europa.eu/fp7/ict/security/engineering_en.html

Please contact:

Thomas Skordas
European Commission
Information Society and Media Directorate General
E-mail: thomas.skordas@ec.europa.eu

EC-ERCIM Strategic Seminar on ICT Security: "Engineering Secure Complex Software Systems and Services", Brussels 16 October 2008

ERCIM and the European Commission are jointly organising a Strategic Seminar on "Engineering Secure Complex Software Systems and Services" on 16 October in Brussels. The seminar is the result of a joint effort of ERCIM, its Security and Trust Management Working Group, and the European Commission (Unit F5 "Security" of DG INFSO).

Scope and Objectives

Responding to the need for closer ties with the European ICT industry, the ERCIM Board of Directors embarked on the initiative to organize a series of annual strategic seminars on current topics within ICT and Mathematics. The rationale behind this strategic decision is manifold: to enhance the impact of research taking place within ERCIM institutes and Working Groups by actively disseminating results towards industrial stakeholders; to expose researchers to ongoing research activities with an industrial take-up potential; and to help bridge the gap between research and industrial practice.

In particular, this ERCIM strategic seminar aims at collecting the relevant academic and industrial expertise in secure software engineering and at linking it with industry's best practices in the field in order to increase the trustworthiness of the resulting ICT systems.

Indeed, the growing complexity of ICT systems and the services they provide creates

demands for a continuously increasing level of assurance on their expected functional behaviour as well as on non-functional properties such as performance, reliability, scalability and in particular security. Today however, the task of secure engineering (from collecting requirements to implementation and operation) of such systems and services is difficult, due to a number of reasons, such as:

- the lack of effective support in writing secure code sections, developing secure systems and assessing their security status
- the lack of adequate methodological support for the elicitation and specification of system-level security requirements based on domain- and application-specific risk analysis
- the lack of support to compare different system implementations with regard to their security properties and expected behaviour.

In order to contribute to addressing these relevant issues, this EC-ERCIM Strategic Seminar will:

- present latest progress on key research and development initiatives in engineering secure complex software systems and services and in achieving ICT system-level assurance
- encourage the dialogue between scientists and industrial players from the field with a view to promoting collaboration; in particular, discuss the balance between rigor-

ous scientific approaches aiming at achieving provably secure systems and cost-benefit considerations

- identify future key research challenges to be addressed in the field.

Interested researchers, practitioners and representatives of stakeholders in this area are invited to attend this event and contribute to its success.

Links:

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

ERCIM WG on Security and Trust Management:
<http://www.iit.cnr.it/STM-WG/>
European Commission's DG INFSO Unit F5 "Security":
http://cordis.europa.eu/fp7/ict/security/home_en.html

Please contact:

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

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

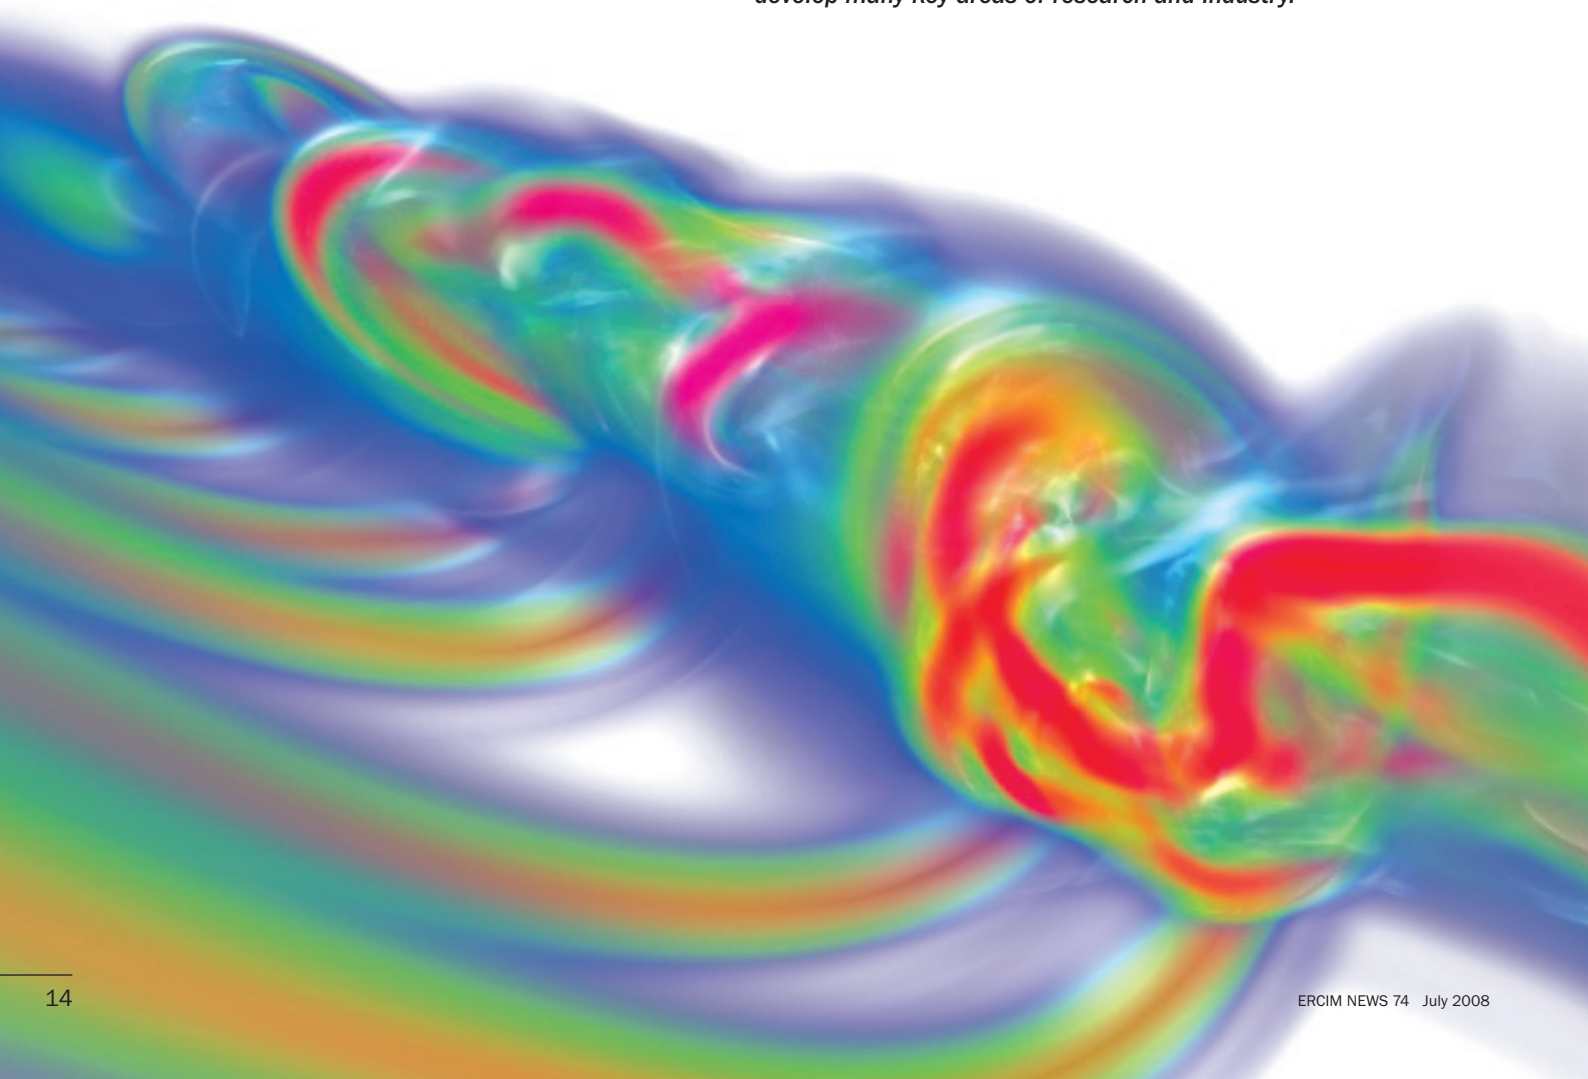
Thomas Skordas,
European Commission
E-mail: Thomas.Skordas@ec.europa.eu

Introduction to the Special Theme

Supercomputing at Work

by Alessandro Curioni and Ray Walshe

Recently, when speaking at Teratec, the conference on High Performance Computing (HPC) in Paris, the European commissioner for the Information Society, Viviane Reding said "Supercomputers are the 'cathedrals' of modern science, essential tools to push forward the frontiers of research at the service of Europe's prosperity and growth". This statement underlines a strategic emphasis being placed on High Performance Computing across the European Union and investment in selected HPC centres aims to develop many key areas of research and industry.



The steady increase in HPC research, the development of more scalable, flexible and energy efficient supercomputer systems in the last 10 years has extended the reach and accuracy of mature supercomputing applications and enabled a myriad of emerging fascinating ones in meteorology, biology, nanotechnology, finance, population dynamics and fluid dynamics to name a few. Supercomputing has now become an indispensable tool to the solution of complex problems in almost any sector.

In this special theme of "Supercomputing at Work" we will see invited and special articles covering diverse research areas. HPC (or Supercomputing) is now widely used every day in climate modelling and weather prediction for our daily forecasts, while members of the European Centre for Medium-Range Weather Forecasting (ECMWF) have systems that predict the expected temperature changes in degrees Celsius as far forward as 2060. On the opposite side of this environmental coin we have reached "Peak Oil" and new energy finds or new sources are required. New research tools like quantitative seismic analysis can be used to detect the presence of hydrocarbons below the sea floor and recover the fuel using CO₂ injection techniques (which is also a means of CO₂ sequestration).

In the fields of electronics and information communication technology, as expected, diverse applied projects have emerged from hardware applications such as nanoCMOS device simulations, exploitation of games architectures in Cell Multi-Processor Arrays and Grid-scheduling systems for high-performance computing applications.

Medicine and computational biosystems (mathematical and computational approaches to large-scale bioinformatics and systems biology) are two of the main biological areas that have shown rapid growth in utilising HPC systems and infra-

structures. Companies are looking at personalized medicine and EU research is progressing in this area including the use of desktop Grids to store e-Health data. Researchers can now use mathematical techniques in the analysis of human bone structures to improve diagnosis of osteoporosis and large-scale computing resources have been applied to cerebrovascular modelling, surgical planning and intervention.

Some really exciting work is also now happening where the large-scale computing architectures allow detailed modelling of complex systems. In this issue we see large-scale immune system models and visualization applications in the field of computational biology. HPC techniques can model the population dynamics of bacterial colonies by taking into account the unique characteristics of each individual bacterial cell and its local environmental conditions. The Proudman Oceanographic Laboratory Coastal Ocean Modelling System (POLCOMS) is prime example of a simulation system that can handle the global marine environment. At its core, POLCOMS predicts the effects of pressure, density etc to provide the forces that drive the flow of water in our seas. The advances in software and hardware have provided the infrastructure for HPC and the strategic investment in science and technology now allows us to study problems of our society including energy, health and climate change.

Complex mathematical and computational techniques that utilise HPC resources now provide solutions to real problems in the many areas such as biological, structural and chemical engineering, fluid dynamics and also thermodynamics. This renewed interest and investment in "Supercomputing" will see the EU competing strongly with the US and Japan within the next decade and hopefully provide some assistance in shaping systems for health, energy and climate for the generations to come.

Please contact:

Alessandro Curioni
IBM Zurich Research Laboratory, Switzerland
E-mail: cur@zurich.ibm.com

Ray Walshe
Centre for Scientific Computing & Complex Systems
(SCI-SYM), Dublin City University, Ireland
E-Mail: Ray.Walshe@dcu.ie
<http://sci-sym.computing.dcu.ie/>

Quantitative Seismic Monitoring Methods

by Martin Landrø

Massive computing power has made possible time-lapse seismic monitoring methods that have been successfully employed in evaluating hydrocarbon fields, particularly in the North Sea.

Over the last two decades there has been a dramatic increase in the monitoring of undersea seismic variation, in part because the technique has been successfully applied to many large hydrocarbon fields, especially in the North Sea. But just one seismic survey can generate 7 TB of data or more, requiring massive computer clusters capable of handing these large data volumes within reasonably short time frames. Additionally, a critical factor in these types of time-lapse seismic survey is the repetition of the seismic experiment between two or more surveys as accurately as possible. Currently, repeated seismic processing can typically be done within a week or less. This approach can be used to improve oil field production and to

monitor the movement of CO₂ injected into undersea reservoirs, whether for storage or for improving oil field performance.

Improved Oil Production

While massive computer processing is an absolutely essential part of seismic variation monitoring, another critical aspect has been the development of high-precision positioning equipment, which is a cornerstone for time-lapse seismic monitoring. In addition to this, the industry has developed dedicated, steerable streamers, as well as cables that are permanently installed in the seafloor above the reservoir. The Valhall Field is one example of this approach. Nine surveys have been

acquired there since 2003, giving a total volume of seismic data of the order of 70 TB.

The careful processing of time-lapse seismic data results in clear images of production-related changes. Figure 1 shows two stacked traces from a given location at the Gullfaks Field in the North Sea, one acquired prior to production starting (1985) and the other after nine years of production (1995). Notice that the two signals are similar from the top down to midway between top reservoir (marked by a red dashed line) and the base of the reservoir (green dashed line). From mid-reservoir down to the base reservoir, represented by the original oil-water contact

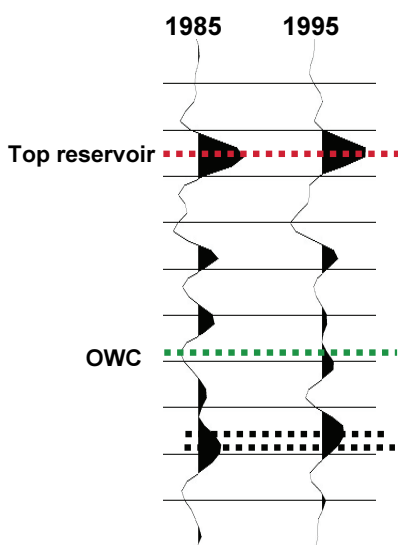


Figure 1: The basic principle of 4D seismic: the repeated seismic signal (recording time along the vertical axis) is similar to the pre-production signal from the top reservoir event (dashed red line) and downwards to midway between the top reservoir and the oil-water contact (OWC, green dashed line). The oil-water contact shows a dramatic change in the seismic response, interpreted as water replacing oil from this level and approximately 40m above. The two black dashed lines below the reservoir show a travel time change (time shift) that is used to compute how much water has replaced the oil.




Figure 2: Long core (left) showing the location for the X-ray cross section (red arrow). Water injection is 50 g/L. X-ray density maps (right) of a core slice showing six time steps during the injection process (from Marsala and Landrø, EAGE extended abstracts, 2005).




Figure 3: Time-lapse seismic data showing monitoring of the CO₂ injection at Sleipner. The strong amplitude increase (shown in blue) is interpreted as a thin CO₂ layer. The dashed red lines indicate top and base of the Utsira sand layer (printed with permission from StatoilHydro).

(OWC), we observe significant changes in the seismic signal, which is interpreted as water having replaced oil in the lower part of the reservoir. Since the sound velocity of water is higher than that of oil, the travel time for a seismic wave passing the whole reservoir section will decrease when water replaces oil, which is shown here as a time shift for the event below the base reservoir (marked by two black dashed lines). By exploiting this 4D seismic monitoring technique at the Gullfaks Field, StatoilHydro has reported a net increase in income of approximately 1 billion USD.

CO₂ Injection and Storage

Interest in CO₂ injection, both for storage and as a tertiary recovery method for increased hydrocarbon production, has grown significantly over the last decade. StatoilHydro has stored approximately 10 million tonnes of CO₂ in the Utsira formation at the Sleipner Field, and several similar projects are now being launched world-

wide. At NTNU our focus has been to develop geophysical methods to monitor the CO₂ injection process, and particularly to try to directly quantify the volume injected from geophysical data.

One way to improve our understanding of how the CO₂ flows in a porous rock is to perform small-scale flooding experiments on long core samples. An example of such a flooding experiment and corresponding X-ray images for various flooding patterns is shown in Figure 2. By measuring acoustic velocities as the flooding experiment is conducted, these experiments can be used to establish a link between pore-scale CO₂ injection and time-lapse seismic data on the field scale.

Figure 3 shows an example of CO₂ influencing seismic data over time. By combining 4D travel time and amplitude changes, we have developed methods to estimate the thickness of CO₂ layers, which makes it possible to estimate volumes. Another key parameter

is CO₂ saturation, which can be estimated using rock physics measurements and models. Although the precision in both 4D seismic methods and rock physics is increasing, there is no doubt that precise estimates are hard to achieve, and we therefore need to improve and combine existing methods in order to decrease the associated uncertainties.

StatoilHydro is acknowledged for permission to use and present the seismic data examples. The Research Council of Norway is acknowledged for financial support for the ROSE (ROck SEismic) project at NTNU.

Link:

<http://www.ipt.ntnu.no/rose/>

Please contact:

Martin Landrø

Norwegian University of Science and Technology, Norway

Tel: +47 7359 4973

E-mail: martin.landro@ntnu.no

Studying CO₂ Sequestration with the Power of Supercomputing

by Klaus Johannsen, Andreas Kopp, Olli Tourunen and Josva Kleist

Computational problems related to the safe storage of CO₂ have become a major focus in environmental science. The computational power required to solve these problems is in principle available at various computing centres. The CO₂ Community Grid project now enables scientists to start large-scale computations with literally eighteen keystrokes. In this way, it brings the power of computing centres to the desks of researchers.

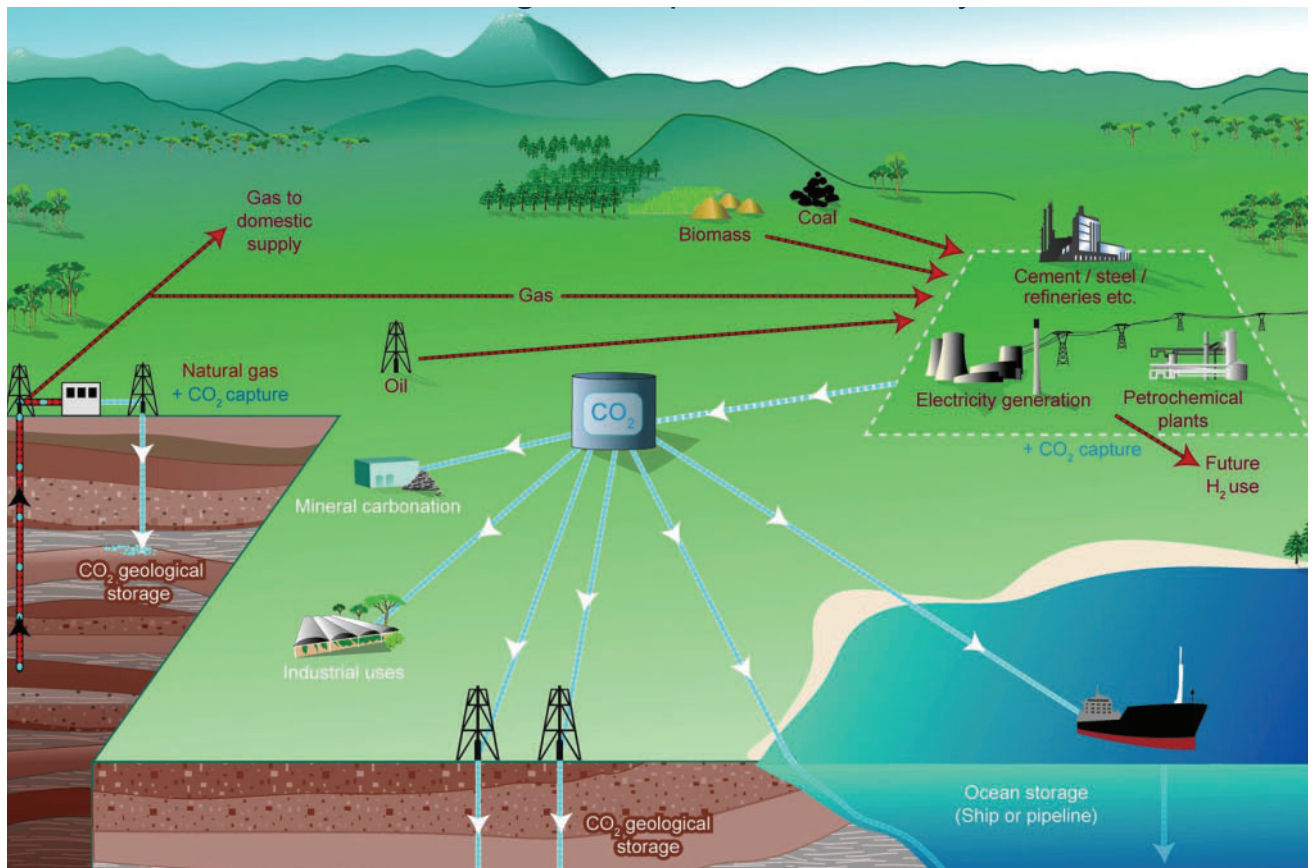
The CO₂ Community Grid (CO₂-CG) project is developing an environment for scientists studying CO₂ sequestration. The project-operated infrastructure provides unified access to distributed supercomputers and supports the parallel execution of applications on the base of MUFTE-UG (Multiphase Flow, Transport and Energy Model - Unstructure Grid), a simulation platform for multi-phase multi-component flow simulations. The project is led by Klaus Johannsen (Bergen Center for Computational Science, University of Bergen) and supported by the Nordic Data Grid Facility (NDGF). Other participants include the Institut für Wasserbau at the University of Stuttgart and the Department of Environmental Engineering at the Technical University of Denmark.

Storage of CO₂ in deep geological formations is regarded as a promising option for mitigating climate changes caused by anthropogenic greenhouse gases. Many countries are currently investigating different aspects related to this topic. Among other approaches, numerical investigation has proven to be an indispensable tool with which to understand the underlying processes.

The first study supported by the project deals with numerical simulations related to the sequestration of CO₂ in deep aquifers. Risk assessments for sequestration scenarios can estimate factors such as leakage rates from aquifers used for sequestration. To determine the critical attributes for potential reservoirs, a parameter study

is conducted using a large number of simulations. Typically each simulation uses 64 processor cores and takes between 4 and 48 hours to complete.

How does it work in practice? Scientists access a service computer that acts as a gateway to various computing centres. They prepare the set of simulation inputs and if necessary modify the source code of the application. After a local test run, they can use a single command to run the simulation on a remote Grid resource (computing centre). The CO₂-CG environment takes care of packing the necessary input files and source code, shipping them to a suitable resource, compiling and running them, keeping track of the progress



Schematic diagram of possible CO₂ Capture and Storage Systems.

Source: Intergovernmental Panel on Climate Change (<http://www.ipcc.ch>) Special Report on CO₂ Capture and Storage.

and returning the results to a user-defined directory.

The heart of the system is a relational database for storing the information necessary for the jobs. The command-line interface inserts all the necessary parameters for the job into the database, and a separate process then performs Grid submission, execution and progress tracking. Stock NorduGrid Advanced Resource Connector (ARC) middleware is used for Grid access, and ARC application environments called runtime environments provide the platform-specific means to run the simulation on each Grid resource (computing centre). Authentication and authorization is based on an industry-standard X.509 security framework with a dedicated Virtual Organization Management System (VOMS) for managing accounts in the user community.

Results so Far

In the first study, scientists at the University of Stuttgart, Germany, investigated the sensitivity of CO₂ distribution in a deep aquifer with respect to various physical parameters. This method is necessary to estimate the risks related to the storage of the greenhouse gas. Within two months the group was able

to perform numerical studies, which used approximately 300,000 CPU hours, or about 15% of the resources allocated to the project for 2008. Without the Grid facilities these numerical studies would have taken about three years of execution time.

The Future

A second research group at the University of Bergen, Norway, is now starting to take advantage of these computing facilities. Apart from ease of use, the Grid community will enhance the existing collaboration between groups and will greatly simplify technical exchanges. As international collaborative research becomes both more common and more important, such possibilities for scientific exchange are essential.

The next steps are to enlarge the CO₂ community by attracting more research groups and further computing resources. In the longer term, the CO₂ Community Grid will be turned into a High-Performance Scientific Computing Grid Community (HPSC-GC) by incorporating other high-performance computing software applications and related scientific groups. In this way the Grid infrastructure will contribute to the broadening of international interdisciplinary research.

Link:

http://www.ndgf.org/ndgfweb/co2_community_grid.html

Please contact:

Josva Kleist
Aalborg University and Nordic DataGrid Facility, Denmark
E-mail: co2@ndgf.org

Numerical Weather Forecasting for Poland

by Maciej Szpindler and Maciej Cytowski

Numerical weather forecasting represents one of the great mathematical modelling challenges applied directly to real-life processes. Methods used for solving corresponding mathematical systems are usually very well suited to work on parallel vector computers. The Interdisciplinary Centre for Mathematical and Computational Modelling (ICM) provides a set of free accessible weather services dedicated to the Polish and wider European communities. Recently ICM opened a new weather service based on the British Unified Model and implemented on an extremely high-resolution grid. In order to achieve a fully operational service we had to optimize the application to fit a specific supercomputer architecture. The result was a success, with an overall speedup of almost three being achieved over the initial implementation.

The weather prediction projects conducted by ICM have quite a large impact on the meteorological community in Poland. ICM serves the public by forecasting the evolution of the state of the atmosphere over central Europe. The results of these computations are available to the general public in the form of detailed weather charts and diagrams presented on ICM's online weather pages. The operational results of the system are used by scientists developing other models, especially oceanographic ones in which forcing from atmospheric winds is treated as an initial condition.

The new generation of numerical weather forecasting models has arisen from continuous development and is the domain of professional meteorological agencies. In particular, the Unified Model (UM) comes from the British Meteorological Office, and COAMPS (Coupled Ocean/Atmosphere Mesoscale Prediction System) was developed by the US Naval Research Laboratory. These models are non-hydrostatic, meaning they use fully three-dimensional equations for




Figure 1: The weather prediction area of ICM's UM v.6.1 model.

atmospheric dynamics. UM and COAMPS also describe new scales of phenomena based on a very dense computational Grid.

Switching these models into operational service means running them permanently using the same daily procedure. This includes transmission of the input data, analysis of initial conditions, model runs and results processing. Operational runs need to produce up-to-

date forecasts and so must be performed as fast as possible. The huge amount of data and computations motivates the use of supercomputers. 'Super' indicates computers that have a lot of memory and fast processor(s), allowing the forecast to be computed in a short time interval.

ICM operates a computational centre that is made freely available to the scientific community. As numerical




Figure 2: The ferry MF Scandinavia and one of the wind charts used by Captain Wojciech Sobkowiak.

weather forecasting has always been a challenging computational task, vector supercomputers were used. In fact, a long line of successive generations of Cray vector machines has seen service at the centre.

Numerical models of atmospheric dynamics are based on a number of differential equations. These are solved on a discrete grid of points that corresponds to a number of geographical positions and vertical levels. This forms a full three dimensional model. The new models are constructed on a spatial grid with 4km resolution or better. Each equation needs initial data and boundary conditions to be solved. These are carefully prepared with use of multi-source observational data.

Currently ICM runs three operational weather prediction models: UM v.4.5 (since 1997), COAMPS (since 2001) and UM v.6.1 (since 2008). In addition the sea-wave prediction model WAM for the Baltic Sea is computed with the use of COAMPS wind fields. Numerical weather prediction at ICM has always been strongly connected to its supercomputing capabilities. On the other hand, ICM staff have always been involved in the optimization of weather prediction computer codes on corresponding supercomputer architectures. The older version of UM was ported and optimized to work on Cray SV1ex. Both the

COAMPS model and UM v. 6.1 were ported and optimized for Cray X1e.

Code Optimization

The ICM's brand new weather prediction system based on UM v.6.1 has two very important novel features. First of all in comparison to previous editions it has some milestone modelling changes, referred to by the Met Office as 'New Dynamics'. Secondly we decided to significantly decrease the grid size to approximately 4km while operating on the same geographic area. All of this made the initial installation of UM v.6.1 very time-consuming, and a single forecast run took more than five hours. Since we plan to perform four operational UM runs per day, we have started to look into opportunities for optimizing the code. The first class of optimizations was vectorization and loop unrolling, which was done with the use of specialized directions for the Cray Fortran compiler. Second, we found several opportunities to substitute some code fragments with machine intrinsic calls. The crucial considerations were those connected with the parallel configuration of the model. The computational domain was divided into a number of long rectangular stripes rather than into square elements. Due to better utilization of the vector processor unit, this resulted in improved performance and an overall speedup of almost three over the initial implementation.

Success Stories

ICM's numerical weather prediction systems have proven their usefulness both for operational and application purposes. Our weather pages are visited frequently by a variety of users, especially those interested in wind prediction, such as wind farm operators, motogliders, windsurfers and sailors. Recently we reported on a spectacular success achieved by Captain Wojciech Sobkowiak, in charge of the ferry MF Scandinavia operating on the Baltic Sea. Captain Sobkowiak used ICM's wind prediction charts to plan a cruise from Nynashamn to Gdansk on 25th of January this year. As he described it, "There was a storm behind us and no wind ahead".

Links:

ICM Numerical Weather Prediction: <http://meteo.pl>
 ICM: <http://www.icm.edu.pl>
 Met Office: <http://www.metoffice.gov.uk>
 NRL: <http://www.nrlmry.navy.mil>

Please contact:

Maciej Szpindler, Maciej Cytowski
 Interdisciplinary Centre for Mathematical and Computational Modelling (ICM),
 University of Warsaw, Poland
 E-mail: m.szpindler@icm.edu.pl,
m.cytowski@icm.edu.pl

Supercomputing: Weathering a Changing Environment

by Ray Mc Grath

Strategically, the use of the Irish Centre for High End Computing (ICHEC) cluster for weather prediction is a more efficient use of resources, and in a very small but satisfying way to date the collaboration helps to reduce the Irish nation's carbon footprint.

Weather forecasting has always had an insatiable appetite for computer resources, driven by a need to improve the realism of numerical models by refining computational Grids or expanding the breadth of physical processes handled by the models. However, scientific aspirations are inevitably tempered by economic reality. For example, Met Éireann, the Irish National Weather Service, launched its operational numerical weather predic-

tion system many years ago by purchasing a very modest supercomputer to support the service. The hardware quickly became dated, and while replacement systems followed every few years, economic restrictions ensured that they never quite matched the cutting-edge status of the original number cruncher. Membership of the European Centre for Medium-Range Weather Forecasting (ECMWF) did eventually bring us back into the super-

computer user club, but only for weather and climate research; for operational forecasting the policy was to depend on in-house systems that were the equivalent of powerful workstations. The issue of climate change overturned this policy.

In 2003 a regional climate modelling project was set up in Met Éireann and it quickly became clear that climate modelling, the 'big brother' of weather fore-

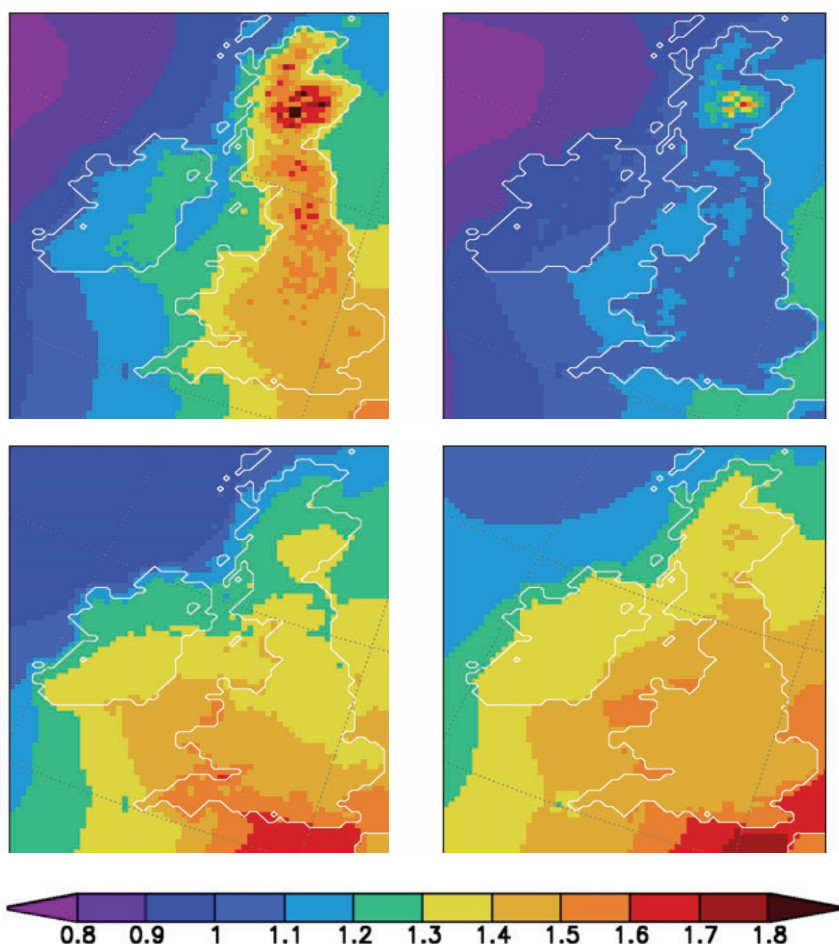


Figure 1: Forecasts of the future climate. Expected temperature change (°C) between the periods 2021-2060 and 1961-2000 for winter, spring, summer and autumn (top left, top right, bottom left, bottom right) for Ireland/UK, based on the mean of eight climate simulations. Only super-computers can deliver this level of detail.

casting, would require substantial computer resources if the project were to be successful. ECMWF's high-performance computing facilities were available through national allocation or Special Project status but allotted quotas were quickly used up. Fortunately, another source of computing power arrived in 2005: the Irish Centre for High-End Computing (ICHEC), a national super-computer centre set up to support Irish researchers. In spite of its relative newness, ICHEC quickly established itself as a reliable resource with excellent support staff. A considerable amount of climate modelling and model development work was carried out on its computers.

This experience prompted Met Éireann to examine its operational computer needs, and eventually led to scientific collaboration with ICHEC in the areas of weather forecasting and climate mod-

elling. As part of this collaboration ICHEC provides computational facilities and support to Met Éireann to enable it to run its operational forecast models. Initially, there were doubts as to whether this collaboration could succeed; Met Éireann's operational production system has stringent delivery and reliability requirements. Eventually, all concerns were addressed in a detailed servicelevel agreement that precisely defines the responsibilities of the partners.

For Met Éireann there are obvious benefits in the collaboration: it provides access to supercomputing resources without having to purchase a dedicated computer or provide infrastructural support systems; it also provides a development path for the operational forecasting system in line with ICHEC's future plans. Strategically, the

arrangement is an efficient use of national resources and, in a very small but satisfying way (given our involvement with climate change issues) it helps to reduce the Irish nation's carbon footprint. To date the collaboration has been very successful.

A big attraction for both sides was the emphasis not only on service provision but on scientific collaboration. ICHEC now joins Met Éireann in contributing to the international HIRLAM (High Resolution Limited Area Model) project, which is dedicated to the development of a new state-of-the-art numerical short-range weather forecasting system for operational use. The large computing facilities available through ICHEC greatly facilitate model development, as extensive resources are required for realistic testing and debugging of the complex code. ICHEC expertise is also invaluable in optimizing the performance of the model on distributed computer systems.

The collaboration also expands Ireland's involvement in climate research, particularly in the climate modelling area. In late 2004 Met Éireann became a partner in the international ENSEMBLES project, a project funded by the European Union under the 6th Framework Programme to assess the future impacts of climate change over Europe. Simulations of the future European climate have been carried out using ICHEC and ECMWF resources. In 2007 Met Éireann and ICHEC became partners in the EC-EARTH project, an international effort to develop a new European Earth System Model based on a global coupled atmosphere-ocean system, with dynamic vegetation, atmospheric chemistry, carbon cycle and ocean biogeochemistry components. Climate simulations with this model will be run at ICHEC.

Links:

- <http://www.met.ie>
- <http://www.ichec.ie/>
- <http://ensembles-eu.metoffice.com/>

Please contact:

Ray Mc Grath
 Head, R&A Division Met Eireann,
 Ireland
 Tel: +353 1 8065 520
 E-mail: ray.mcgrath@met.ie

Supercomputing at Work in the nanoCMOS Electronics Domain

by Richard Sinnott, Campbell Millar and Asen Asenov

The UK Engineering and Physical Sciences Research Council in collaboration with leading electronics design houses, chip manufacturers and tool vendors has given funding of £5.3m to apply e-Science and Grid technology to some of the fundamental challenges caused by the decreasing 'atomistic' scaling of Complementary Metal Oxide Semiconductor (CMOS) transistors. Progressive scaling of CMOS transistors, as tracked by the International Technology Roadmap for Semiconductors and captured in Moore's law, has driven the phenomenal success of the semiconductor industry, delivering larger, faster and cheaper circuits. However the years of 'happy scaling' are over and the semiconductor industry must now address the challenges caused by intrinsic parameter fluctuations stemming from the fundamental discreteness of charge and matter. These parameter fluctuations are truly stochastic and cannot be eliminated by tighter process control.

The nanoCMOS project began in October 2006 and is due to run for four years. It involves collaboration between leading device-modelling and circuit and system design research groups at the Universities of Glasgow (PI Prof Asenov), Edinburgh, Manchester, Southampton and York. This strong academic group is enhanced by robust links with industrial partners including leading tool vendors and design companies such as Synopsys, Freescale, ARM, Fujitsu, National Semiconductor and Wolfson Microelectronics among others.

The increasing device variability of CMOS devices demands revolutionary changes in the design of, and design methods for, integrated circuits and systems. To achieve this, it is now recognized that a real paradigm shift must occur and strong links must be established between system, circuit and fundamental device technology research and design groups to allow modern integrated circuits to cope with the individual behaviour of statistically varying transistors on a chip. In the past, the assumption has been that there is a single idealized transistor used across a chip, but as the scale of devices shrinks this no longer holds. Instead, design methods must evolve to accommodate the increasing statistical variability of transistors and the impact this can have on circuit and system performance.

Changing design rules for device architecture and variability adds significant complexity to the design process, requiring the orchestration of a broad spectrum of design tools by geographically distributed teams of device experts and circuit and system designers. In

nanoCMOS this is being addressed by embedding e-Science technology and know-how across all nanoCMOS electronics design groups, revolutionizing the way in which these disparate groups currently work and more importantly, providing new insights for industry on the challenges faced in the nanoCMOS domain and on how the global semiconductor industry can address them.

In order to study the statistical fluctuations introduced by the atomic structure of transistors it is necessary to perform 3D simulations of huge ensembles of hundreds of thousands of devices, rather than a single representative device. Given the increasing number of transistors in modern chips, simulation of very large statistical samples of devices is required to allow statistically




Figure 1: Potential/dopant distribution for a statistically rare device.




Figure 2: Threshold voltage variation as a function of the number of dopant atoms in the transistor.

rare devices with potentially fatal effects on circuit performance to be examined. This demands access to huge distributed high-performance computing resources including the UK e-Science National Grid Service and Scot-Grid, and a wide variety of other resources including Condor pools and campus clusters across partner sites.

However this is not simply another large-scale simulation problem, since the commercially sensitive nature of the problems and stringent IP protection demands fine-grained security on access to, and usage of, licensed software, providing protection for the intellectual property associated with circuit and device designs, data and simulations belonging to industrial partners and key stakeholders. To this end, the project has developed infrastructure that provides security throughout. This includes exploitation of Kerberos for secure global file-based access through the Andrew File System; authorization technology such as PERMIS for definition and enforcement of access policies using centralized attribute authorities such as the Virtual Organization Membership Service (VOMS); and simple user-oriented access to the project por-

tal through the Internet2 Shibboleth technology using the UK Access Management Federation. This portal contains a variety of services for atomistic device simulation, compact model generation and circuit simulation, as well as mechanisms for managing the large and heterogeneous data and metadata associated with these simulations.

The project is now well advanced and the initial prototype infrastructure is ramping up for large-scale scientific usage. Up until now, due to the computational complexity of 3D device simulation, studies of variability have tended to be based on small ensembles of typically up to 200 devices. We have simulated ensembles in excess of 100,000 3D devices for 35nm and 13nm gate-length devices. Ensembles of this magnitude are shedding new light on the effects of atomic structure variation on the behaviour of devices, especially extreme limits of device variability. Furthermore, based on these simulations, we have been able to examine the effect of device variability at a simple circuit level and have simulated over 1 million CMOS inverters using random configurations of devices. Figure 1 shows the potential and dopant position

of a statistically rare device. Figure 2 shows the threshold voltage variation as a function of the number of dopants. The second phase of this project will use the methodologies developed to study larger and more advanced circuits and systems and to further explore the impact of atomistic variability of transistors on the design process. More information on the nanoCMOS project is available on our Web site.

Links:

UK Engineering and Physical Sciences Research Council:

<http://www.epsrc.ac.uk>

The nanoCMOS project:

<http://www.nanocmos.ac.uk>

UK e-Science National Grid Service:

<http://www.ngs.ac.uk>

ScotGrid: <http://www.scotgrid.ac.uk>

Please contact:

Asen Asenov

(science-related questions)

University of Glasgow, UK

E-mail: a.asenov@elec.gla.ac.uk

Richard Sinnott

(e-Infrastructure-related questions)

University of Glasgow, UK

E-mail: r.sinnott@nesc.gla.ac.uk

SPRANKER: A Discovery Tool to Rank Service Providers Using Quality of Experience

by Domenico Laforenza, Franco Maria Nardini, Fabrizio Silvestri and Gabriele Tolomei

The HPC-Lab at ISTI-CNR is investigating the topic of service discovery with the aim of supporting service-oriented architectures (SOAs) on Grid computing infrastructures. We present a brief overview of SPRanker (Service Provider Ranker), a discovery tool that, unlike the usual service discovery components, retrieves provider information rather than service descriptions. At its core SPRanker exploits a score formula based on information retrieval that takes into account judgments expressed collaboratively by past service users.

The HPC Lab at ISTI-CNR has developed a service provider ranker (SPRanker) tool. The SPRanker module intervenes between the three main service-oriented architecture stakeholders: 'providers' that publish services, 'users' that discover and bind services, and 'brokers' that act as a provider's medium for spreading information on services to users. Users, instead, use a broker to locate and select the services they need. Figure 1 depicts the publish-discover-bind process that typically takes place in real-world Grid-service-based SOAs.

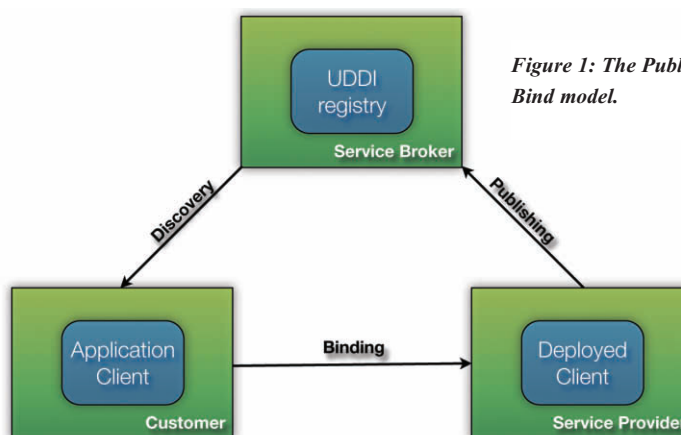


Figure 1: The Publish-Discover-Bind model.

Modern approaches to quality of service (QoS) within SOAs adopt service-level agreements (SLAs) as a way of defining constraints that must be satisfied by customers and providers. In this respect, SLAs represent a 'best effort' strategy for selecting service providers. Simply put, in this kind of system the broker is usually a UDDI (Universal Description, Discovery and Integration) registry publishing WSDL (Web Services Description Language) of the stored services. Providers publish a service by pushing its description into the registry. Customers discover services by querying the registry for a service URI. Finally, the binding phase consists of invoking the actual service through Simple Object Access Protocol (SOAP).

We hypothesize that in the near future, the world will be populated by thousands of millions of services from different providers. Like the Web today, where the same information is supplied by different Web sites, many different providers will supply diverse services in the future Net.

In order to develop a reliable, scalable, highly available and highly performing service, it is necessary that the discovery phase provide the best possible set of matching services (ie that it have a high level of precision). In SLA-based SOAs, once a service has been found it is bound to the client only if the SLA-template the provider offers is appropriate for the customer.

SPRanker not only returns a flat list of results but also ranks the various providers according to some quality metric. The service designer eventually chooses the provider from the ranked list. Note that SPRanker is different from UDDI registries, which are not capable of retrieval on the basis of QoS information. The use of UDDI corresponds to composing workflows according to a 'best effort' strategy.

Our Contribution

Our ranked discovery service implements a novel ranking schema based on solid information retrieval theory, namely the vector space model, by considering historical information on expired SLAs. The ranking score is in fact based on the assumption that provider performance (in terms of QoS) should be evaluated collaboratively by considering user feedback.

The vector space model represents an object as a vector in \mathcal{R}^n where each dimension corresponds to a separate term. If a term occurs in the object, its value in the corresponding vector entry is non-zero. SLAs are the objects that are modelled as vectors in \mathcal{R}^n . Here, n is the number of possible SLA terms, and each SLA term is mapped into a particular dimension. To keep the model as simple as possible we consider only unit vectors. The normalization is done in such a way that all the

weighted by the time elapsed since the SLA was issued. Presenting a Query-SLA, SPRanker seeks a list of providers offering a particular service ordered by similarity with the query.

The architecture of SPRanker is composed of three modules: gatherer, indexer and query server. The gatherer collects data from (positively) expired SLAs. We only consider positively expired SLAs because we want to discriminate between good and bad service




Figure 2:
Architecture
of SPRanker.

vector coordinates will range between 0 and $1/n$.

An SLA-vector is defined as a unit vector representing a successfully completed service-level agreement issued to a service provider at a given time. Each value of the vector is the value associated to a term of the SLA template. For example, could represent the SLA of a service provided at time T running on a 2-way SMP, with 1TB free disk space and at the cost of 0.04\$.

Queries in SPRanker are called Query-SLAs. A Query-SLA is a unit vector where each value can assume one of the following values:

1. A reference value for term T_i of the SLA template
2. $^\circ$, meaning that we do not want to take into account the i -th term
3. \bullet , meaning that the i -th term may assume any value between 0 and $1/n$.

Assume a Query-SLA and a set of SLA-vectors representing an SLA successfully issued at time T , by a provider, on a particular service. A similarity function that takes into account the provider, service name and SLA issue time is defined. We define $sim=0$ if either the provider or the service name differs. Otherwise the similarity is defined as the sum of the common terms shared by Query-SLA and an SLA-vector

provisioning, and because we want to avoid satisfied customers incurring false bad judgements from malicious partners (clients or customers) willing to lower a provider's score. The gatherer can act in two different modes – push-based and pull-based.

When in push-based modality, the gatherer receives SLAs directly from providers and customers. In contrast, pull-based mode is used to periodically poll known providers for up-to-date information. The indexer is used to transform SLAs collected by the gatherer into a machine-readable format. The query server has been implemented as a Web service. It offers two distinct methods, one for each kind of query answered by SPRanker.

This research has been carried out in the frame of S-Cube, the "European Network of Excellence in Software Services and Systems" funded by the European Commission's Seventh Framework Programme.

Link:

<http://www.s-cube-network.eu/>

Please contact:

Franco Maria Nardini

ISTI-CNR, Italy

Tel: +39 0050 3153055

E-mail: francomaria.nardini@isti.cnr.it

Design of a Functional Architecture for the Management of Cluster Resources and Services through the Web

by Juan Antonio Ortega, Jorge Cantón, Ana Silva, David Bosque and Francisco Velasco

In 2006 Junta de Andalucía created the Andalusian Supercomputing Network (RASCI). RASCI consists of supercomputing nodes distributed geographically throughout Andalusia that provide the region with a large number of computing resources. Increased network bandwidth, more powerful computers and acceptance of the Internet have driven a growth in demand for new and better ways to utilize high-performance technical computing (HPTC) resources.

The Andalusian Scientific Computing Centre (CICA) is a supercomputing centre in Andalusia that provides supercomputing nodes and the Andalusian e-Science Portal to RASCI (see Figure 1).

CICA Infrastructure

Currently, CICA has a distributed cluster. It consists of a set of nodes with Intel Core 2 Duo or Intel Quad Core CPUs and 4 GB or 8 GB of RAM, connected through gigabit ethernet. The batch systems installed are the Sun Grid Engine (SGE) and Condor interface. At present, the Web interface application allows users to queue their jobs into Condor and SGE for subsequent execution by the cluster.

CICA promotes the use of open-source software, and the supercomputing nodes run under Linux, make use of open-source libraries and utilities, and use several OS.

The Andalusian Supercomputing Network will allow nodes located at different universities and research centres to be joined together in a unified computing system. To achieve this, a technical committee (for infrastructure management), an access committee (for access management) and a strategic committee have been established.

Other services are offered by CICA to users of the Andalusian Supercomputing Network, including assistance in parallelizing applications to take advantage of the Grid infrastructure. This involves providing programming courses, expert advice and consultancy for optimization and parallelization of applications and the use of message-passing libraries for the efficient use of distributed memory. With these services CICA intends to increase the number of users who can potentially take advantage of the super-

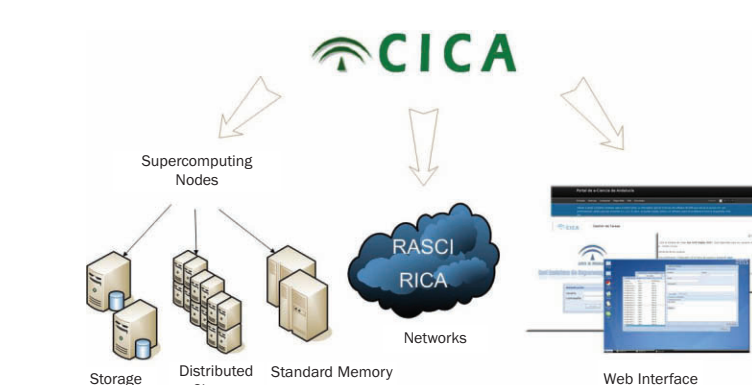


Figure 1: CICA connections.

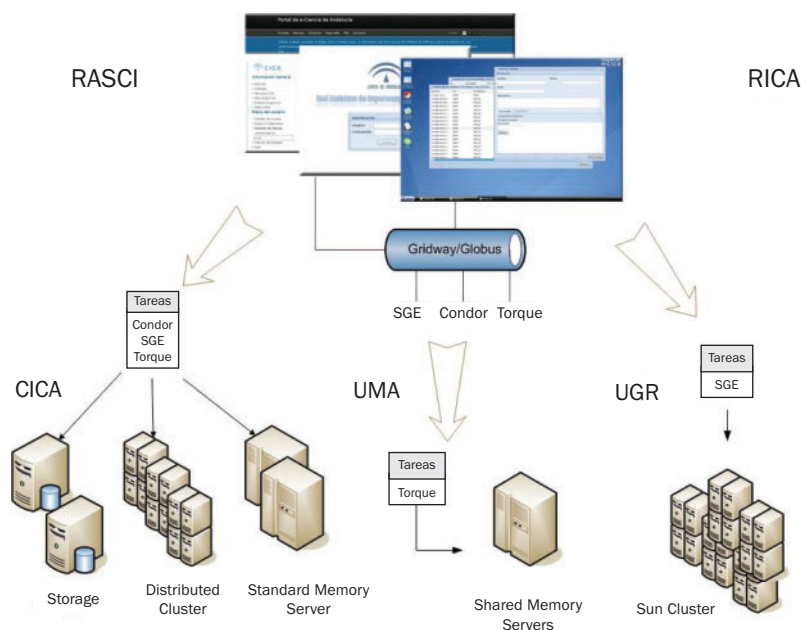


Figure 2: Supercomputing infrastructure.

computing infrastructure, rather than limiting its use to those who are already technically proficient (see Figure 2).

All of this makes it possible to bring supercomputing technology to a heterogeneous community, allowing for faster technological progress, and also to increase the demand for this type of service, due to its simplicity of use.

Web Interface

Rather than being graphical or accessible through a Web interface, cluster queue manager interfaces tend to be accessed through a command line interface from a Linux or Windows console.

Different interfaces have been developed for the Condor queue manager, each with a specific purpose depending

on who is making use of it. One of the services to which this queue manager's developers have contributed is the use of Web services (SOAP-WSDL). For its implementation they have developed the BirdBath package, which, by means of WSDL calls, allows jobs in the cluster to be submitted and managed. The main inconvenience presented by these calls is the ease of discovery and indexation by search engine Web crawlers, since they are accessible through URLs. They can thus be found and stored through search engines. Moreover, the management capabilities are limited to the functionalities implemented by BirdBath developers.

The architecture presented by the CICA-named Cluster Looking Class (CLG) allows the user to make use of resources through the Web, eliminating the need to enter commands through a command line console. This ease of use

extends the reach of computing services to a greater number of users as a technical profile is no longer required. At present, this interface is operational at the Andalusian e-Science Portal, giving registered users access to computational resources connected to the Andalusian Computing scientific Network (RICA) and Andalusian Supercomputing Network (RASCI).

However, with the ability to access the system through the Web, security becomes a critical aspect. In a system such as this, the high level of computing power turns it into a very powerful tool for illicit uses such as brute-force password cracking. In order to avoid this problem, CLG defines both a logical and a physical architecture that allows the user to manage and submit jobs and to receive their results, uncoupling the physical architecture of the cluster from the Web interface and

thus putting an end to security problems.

Further Work

In the future, the supercomputing resources of CICA will grow and will be included in this infrastructure. Our estimation is that this will lead to bigger massive storage, shared memory servers, more distributed nodes and new batch systems.

Links:

Andalusian e-Science portal:

<https://eciencia.cica.es>

Condor Project HTC:

<http://www.cs.wisc.edu/condor>

Please contact:

Juan Antonio Ortega

Centro Informático Científico de Andalucía, Spain

E-mail: ortega@cica.es

Exploitation of Cell Multi-Processor Array in Solution of Spatio-Temporal Dynamics

by Zoltán Nagy, László Kék, Zoltán Kincses, András Kiss and Péter Szolgay

Array computers can be useful in the numerical solution of spatio-temporal problems. IBM has recently introduced a topographic array processor called the Cell Broadband Engine (Cell BE). Researchers at the Cellular Sensory Wave Computers Laboratory of SZTAKI in collaboration with the Department of Information Technology, Pázmány Péter Catholic University, Budapest and the Department of Image Processing and Neurocomputing, Pannon University, Veszprém, have implemented a Cellular Neural Network (CNN) simulation kernel on the Cell BE. The CNN simulation is optimized according to the special requirements of the Cell BE and implements both linear and nonlinear (piecewise linear) templates. We have used the CNN simulation kernel to solve Navier-Stokes partial differential equations (PDEs) on the Cell architecture.

Array processing is a good candidate for increasing computing power by using parallel computation. Additionally, it can help to solve architectural problems (eg distribution of control signals on a chip). The effectiveness of implementations is measured by a set of parameters, namely by the silicon area (A), execution time (T), dissipated power (P), consumed energy (E) and number of input/output pins (#I/O).

A number of different implementations of array processors are commercially available. In this project we have mainly concentrated on the topographic IBM Cell heterogeneous array processor architecture, because its development environment is open source and we

wished to compare the results with those from our previous FPGA (field programmable gate array)-based implementations.

The Cell Broadband Engine Architecture (CBEA) is designed to achieve high-performance computing with better area/performance and power/performance ratios than the conventional multi-core architectures. The CBEA defines a heterogeneous multi-processor architecture where general-purpose processors called Power Processor Elements (PPE), and Single Instruction Multiple Data (SIMD) processors called Synergistic Processor Elements (SPE), are connected via a high-speed on-chip coherent bus called an Element

Interconnect Bus (EIB). The CBEA architecture is flexible and the ratio of the different elements can be defined according to the requirements of different applications. The first implementation of the CBEA is the Cell Broadband Engine (Cell BE or informally Cell) designed for the Sony PlayStation 3 game console, which contains one PPE and eight SPEs.

In this work we have concentrated on an efficient CNN implementation on the Cell architecture. The basic CNN simulation kernel was successfully implemented on the Cell BE, and both linear and nonlinear CNN arrays can be simulated. The kernel was optimized according to the special requirements of


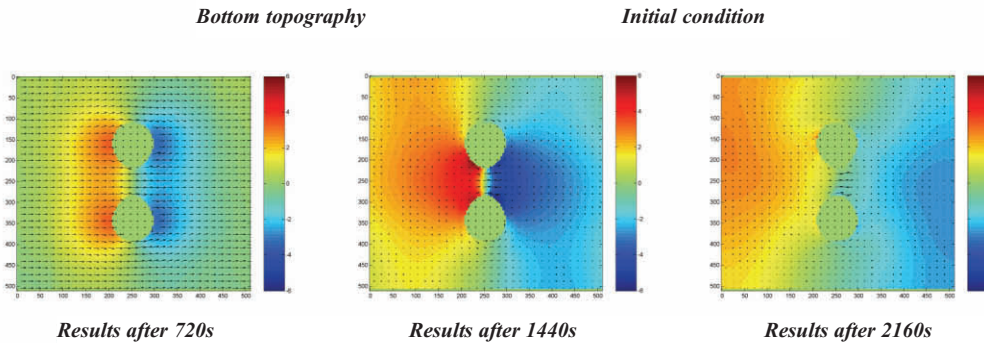


Figure 1:
2D snapshots of ocean current modelling.



the Cell architecture. Performance comparison showed that a roughly sixfold speedup can be achieved over a high-performance microprocessor in the single SPE solution, while the speedup is 35-fold when all eight SPEs are utilized. When using nonlinear templates the performance advantage of the Cell architecture is much higher. In a single SPE configuration, a 64-fold speedup can be achieved, while the use of eight SPEs means the performance is 429-fold improved.

In addition, the CNN paradigm was used to solve a complex spatio-temporal problem. Namely, the 3D Princeton Ocean Model was implemented on Cell BE and

a significant improvement in performance was achieved. Our solution was optimized according to the special requirements of the Cell architecture. Performance comparison showed that an approximately 17-fold improvement can be achieved over a high-performance microprocessor in the single SPE solution, while the speedup is 85-fold when all six SPEs are utilized.

Figure 1 illustrates flow through a channel that includes two islands at the centre of the domain. The size of the modelled ocean is 1024km x 1024km, the north and south boundaries are closed, the east and west boundaries are open, the grid size is 128x128, and the grid resolution

is 8km. The simulation time-step is 6s and 360 iterations are computed. In the future, further speedups might be achieved by using the full power of the Cell architecture on Cell processor-based IBM blades.

Links:

- <http://www.research.ibm.com/journal/rd/494/kahle.html>
- <http://cnn-technology.itk.ppke.hu/>
- <http://www.aos.princeton.edu/WWWUBLIC/htdocs.pom/index.html>

Please contact:

Zoltán Nagy
SZTAKI, Hungary
E-mail: nagy@sztaki.hu

IANOS – Efficient Use of HPC Grid Resources

by Philipp Wieder, Wolfgang Ziegler and Vincent Keller

IANOS - Intelligent ApplicationN-Oriented Scheduling framework is a Grid-scheduling system that provides a generic job-submission framework for automatic optimal positioning and scheduling of high-performance computing applications.

IANOS solves the nontrivial problem of scheduling applications within a Grid of high-performance computing (HPC) machines, through a novel approach of combined cost and execution time evaluation based on job requirements, application and resource characteristics, and historical data. These algorithms are the backbone of a service-based framework that automatically finds the resources best suited to

a certain application request, and negotiates the specific Quality of Service demands with the respective resource providers. For that purpose, IANOS' inherent negotiation framework uses standard service-level agreements, a mechanism suitable for both academic and industrial purposes.

The value of IANOS for users (or customers) is evident: they need not deal

with the peculiarities of the underlying Grid system and its topology, but merely define their job-specific requirements. These are processed by IANOS to schedule the job and to generate the Grid middleware-specific job descriptions. The latter are realized through an adapter customized per middleware, since IANOS in general is middleware-agnostic, a fact that minimizes the overhead of its deployment.

A number of Grid schedulers currently exist, including some that are application-oriented. They serve different needs, user groups and applications, apply different scheduling algorithms, and integrate (or not) with different Grid middlewares or infrastructures. IANOS is such a Grid scheduler, but at the same time is more than that. It comprises a number of components and suggests a certain modus operandi for an application-based Grid scheduling, but the system has been designed to match the requirements of a generic Grid-scheduling framework. IANOS:

- is based on standards wherever possible. It currently implements the Web Services Resource Framework (WSRF), Job Submission Description Language (JSDL), WS-Agreement, and GLUE (Grid Laboratory for a Uniform Environment), and the IANOS team is cooperating closely with some of the standardization groups. In particular the work towards a generic scheduling architecture conducted by Open Grid Forum's 'Grid Scheduling Architecture Research Group' is worth mentioning.
- applies novel scheduling methods for HPC Grids. This is done through the CoreGRID Network of Excellence by close collaboration between application users, system administrators and Grid developers.

- is designed to serve multiple application areas. The architecture is service-oriented, based on Web services, modular and as generic as possible. Workflow scheduling, co-allocation and additional applications are already on the research agenda.
- is middleware-agnostic. The system has been designed so that it can be adapted to different middlewares with minimal effort. Currently a UNICORE (Uniform Interface to Computing Resources) adapter is available and a Globus adapter is under development.
- is targeting production systems. Once the extensive tests have been finalized, IANOS will be used in production mode.

The scientific work done within CoreGRID involves the research groups on Grid Scheduling Architecture, Integration of Intelligent Scheduling Service (ISS) into the VIOLA (Vertically Integrated Optical Testbed for Large Applications)/MetaScheduling Environment, performance prediction using the Gamma model, and service-level agreements for resource management and scheduling. In addition to this, IANOS comprises results from the collaboration with the SwissGrid Association SWING, and collaboration and mutual result exchange with several standardization groups at the Open Grid Forum.

Although the focus is on Grid scheduling, areas like business-oriented Grids, security, information management, accounting and so on are also of great importance. The core components of IANOS are as follows:

(1) The broker of the IANOS middleware uses two models: the cost function model and the execution time evaluation model. These are based on a parameterization of the applications and resources. The cost function model calculates the cost value for each candidate resource. The execution time evaluation model forecasts the execution time of a given application on a given resource, a prediction based on knowledge about the CPU node performance of the applications. The broker itself relies on information coming from an information service and diverse monitoring services.

(2) The MetaScheduling Service (MSS) applies multi-level scheduling strategies with interaction between local resource management systems and higher-level Grid-scheduling entities. Collaboration of Grid-level schedulers with local resource management systems in a heterogeneous environment raises a number of issues, such as determining the level of interaction, coping efficiently with heterogeneity, negotiating with entities that in general do not offer this property, and deciding which negotia-

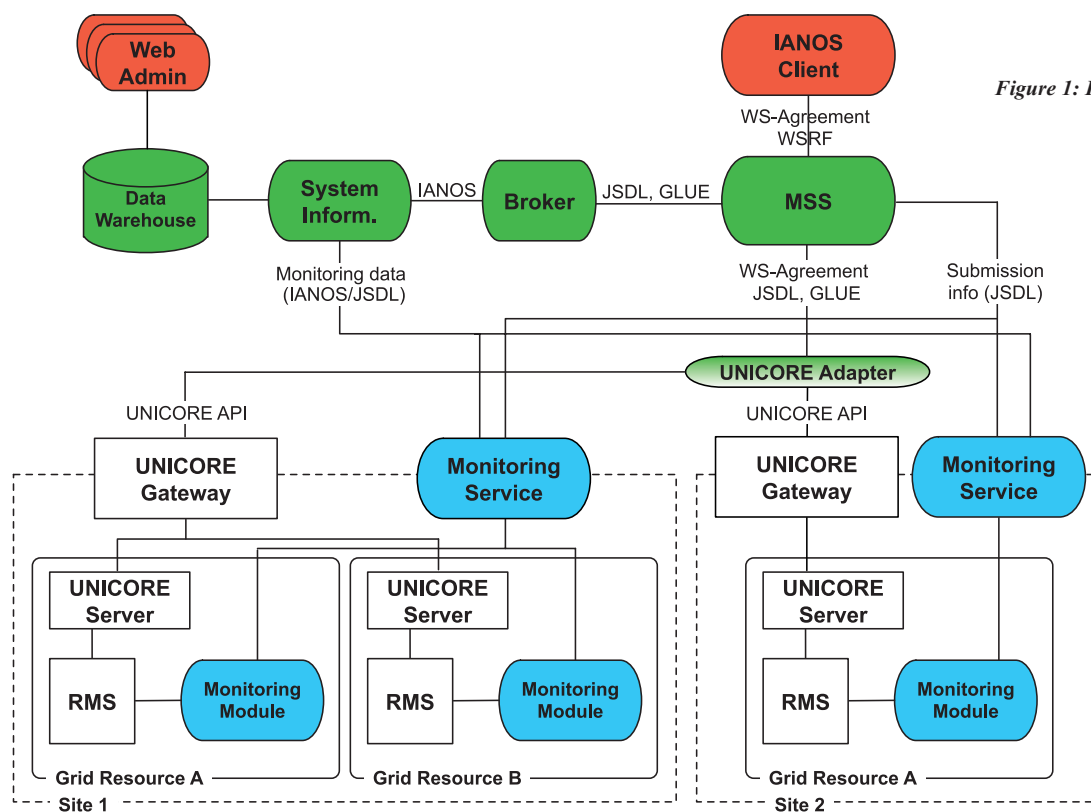


Figure 1: IANOS architecture.

tion protocol to use. Here the IANOS team collaborates closely with the Grid Scheduling Architecture Research Group of the Open Grid Forum.

(3) Service-Level Agreements (SLAs) serve as electronic contracts that formalize the level of service upon which the user and the MetaScheduling Service agree. IANOS implements an SLA framework based on the WS-Agreement standard. Since WS-Agreement also serves as the foundation for a number of business-oriented Grid systems, IANOS may also be applied in a business context.

IANOS has recently been demonstrated at the Open Grid Forum 23 in Barcelona. The demonstration showed the framework in action, spanning three heterogeneous HPC resources at three sites residing in three different adminis-

trative domains: Ecole Polytechnique Fédérale de Lausanne, Fraunhofer SCAI and Dortmund University of Technology. Each site has the necessary IANOS components and applications installed, including:

- MXM for matrix-matrix multiplication. This core component of many scientific applications is dominated by CPU speed requirements.
- MXV for sparse matrix times vector multiplication. MXV is dominated by memory bandwidth.
- SpecuLOOS is a spectral element solver for the Navier-Stokes equations in 3D, which runs in parallel and is dominated by point-to-point communication.
- LBM-Solver is a Lattice-Boltzman Method solver in 3D. Its dominating characteristic is multicast communications.

Links:

<http://www.ianos.org>
<http://www.coregrid.eu>

Please contact:

Philipp Wieder
Dortmund University of Technology,
Germany
Tel: +49 231 755 2767
E-mail: philipp.wieder@udo.edu

Wolfgang Ziegler
Fraunhofer SCAI, Germany
Tel: +49 2241 14 2258
E-mail: wolfgang.ziegler@scai.fraunhofer.de

Vincent Keller
Ecole Polytechnique Fédérale de
Lausanne, Switzerland
Tel: +41 21 693 38 56
E-mail: Vincent.Keller@epfl.ch

Towards Personalized Medicine: High-Performance Computing in the Life Sciences

by Olaf Schenk, Helmar Burkhart and Hema Reddy

Personalized medicine based on high-performance computing (HPC) has the potential to transform health care and to dramatically improve clinical outcomes. This new approach to medicine provides clinical researchers with fast access to information about individual patients. A promising exploitation of the Cell architecture for personalized medicine in biomedical life science applications is the subject of an HPC research project between the University of Basel and IBM.

Understanding the factors that influence human health and cause disease and using that understanding to develop treatments has always been the driving force behind medical research. In personalized medicine, the goal is to perform real-time analysis on a patient (using techniques such as magnetic resonance imaging) that can then be used to plan an individual course of treatment for that patient. Biomedical applications offer significant benefits both to the public and the scientific community, and consequently are emerging as challenging opportunities for innovation at the meeting point between medicine and computer science.

Frontiers in Medicine and HPC

In the world of high-performance computing, all signs point towards a future populated with multi-core heterogeneous processing technology. We are currently witnessing the emergence of multiple-core processors such as Intel

Quad-Cores, graphics processing units and the Cell Broadband Engine™ (Cell BE), which are bringing an unprecedented combination of low cost and low power to high-performance computing. A cooperative effort between the University of Basel and IBM is exploring the use of the Cell BE architecture in the individual biomedical planning of hyperthermia cancer therapy.

Hyperthermia is a promising treatment modality for various types of cancer. The technique involves heating the tumour with electromagnetic fields, generally using antenna arrays to focus the energy. In planning the therapy, the therapeutically optimal antenna parameters for the applicator are determined for each patient and the temperature distribution is predicted by solving the 3D Pennes bio-heat transfer equation. Although this can be a demanding task, a planning tool can greatly help clinical researchers to model and simulate the

medical treatment. In this Cell BE project, we are addressing the practical concerns of fitting into a clinician's standard work-flow, and the question of how to use optimal algorithms and relevant HPC architectures to maximize application performance.

The Cell BE in a Nutshell

Cell BE is a heterogeneous multi-core processor that consists of one 64-bit power processing element (PPE) as the main processor and eight co-processors called synergistic processing elements (SPEs), which act as vector processors specialized for running computationally intensive applications. The latest chip generation of the Cell processor is now able to perform double-precision calculations with a peak near 200 GFlop/s. The SPEs are more specialized processors and consist of a synergistic processing unit (SPU) and a memory flow controller (MFC). An SPU is capable of performing SIMD (Single Instruction,




Figure 1: Simulation of temperature distribution in a human body as influenced by electromagnetic fields.

Multiple Data) floating-point arithmetic and performs best when used in a streaming model. The MFC handles direct-memory access (DMA) transfers between the local SPE memory region of size 256 kB and the main memory of size 32 GB.

One of the radical changes in the Cell BE architecture is the elimination of the cache for the SPEs. The SPEs cannot access the main memory directly, so data required by the SPE must be explicitly transferred to and from the SPE local store by the use of asynchronous DMA transfers, which are handled by the MFC. The memory model, along with this specialization of tasks between SPEs and PPE, is a significant factor in the improvement in performance – approximately an order of magnitude – for key scientific applications and area-and-power efficiency.

Impact on Life Sciences

An order of magnitude improvement in performance over existing CPU technology is a disruptive change that can dramatically alter aspects of personalized medicine. A computational task that previously would have taken a year can now be completed in a few days; hour-long computations suddenly become interactive because they can be completed in seconds with the new Cell technology; and previously intractable real-time processing tasks now become tractable.

Since it is the temperature increase in hyperthermia cancer therapy that causes the increased cell death in tumours, it is of the utmost importance to determine the temperature distribution correctly and with high resolution. For this, the 3D Pennes bio-heat equation with temperature and time-dependent tissue parameters for the blood vessels can be solved interactively on the Cell architecture with 10^8 voxels. Validation of the Cell BE simulation is one of our future targets and will be performed with the Foundation for Research on Information Technologies in Society (IT²S), Zurich, and the hyperthermia unit of the Erasmus Medical Center in Rotterdam.

Links:

<http://fgb.informatik.unibas.ch/research/IBMFacultyAward/index.html>
<http://www-03.ibm.com/technology/cell>

Please contact:

Olaf Schenk
 University of Basel, Switzerland
 E-mail: olaf.schenk@unibas.ch

Hema Reddy
 IBM Corporation
 Systems & Technology Group
 Austin, USA
 E-mail: hemlata@us.ibm.com

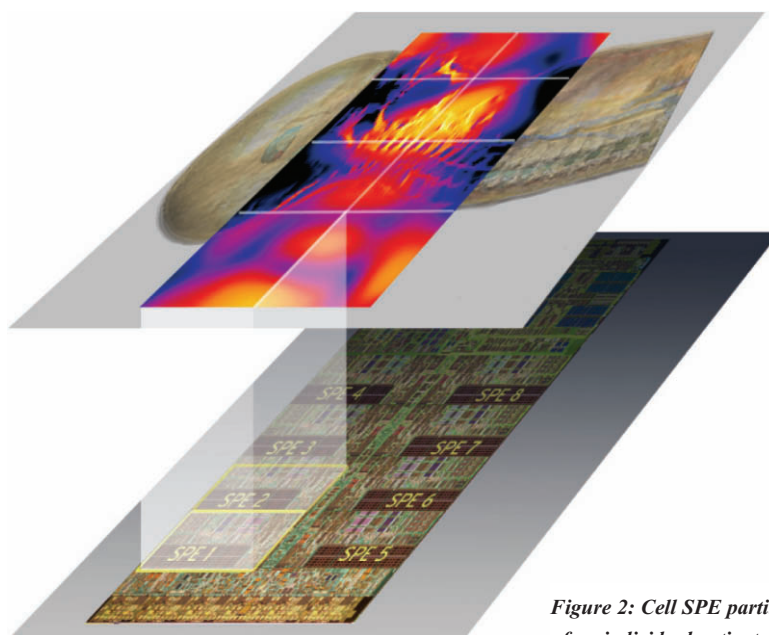


Figure 2: Cell SPE partitioning of an individual patient.

Microstructural Finite Element Analysis of Human Bone Structures

by Peter Arbenz and Ralph Müller

According to the World Health Organization (WHO), the lifetime risk for osteoporotic fractures in women is close to 40%, while in men it is about 13%. Osteoporosis is second only to cardiovascular disease as a leading health-care problem. Therefore it is of paramount importance to be able to predict osteoporosis as early as possible in order to be able to take effective measures to prevent its progression. This project is about locating positions in the bone that are subject to excessive stresses and thus are prone to failure. Supercomputing offers novel approaches to simulating and visualizing larger pieces of bone with unprecedented resolution. Simulations of bones that have been strengthened by implants have also been successfully conducted.

For the clinician, predicting fracture risk for individual patients essentially amounts to the quantitative analysis of bone density. However, several studies have shown that bone strength, an indicator for bone fracture risk, is only moderately predicted by bone density. This is not surprising: bones are not solid structures, but are made up of an outer shell of compact bone enclosing a core of spongy, trabecular bone. This trabecular bone, which is located at the end of long bones, has a porous structure that contributes significantly to the load-bearing capacity of the human skeleton. It is not a random structure; rather its trabeculae are typically aligned with the main loading direction.

The investigation of the mechanical properties of trabecular bone presents a major challenge. This is due to its high porosity and complex architecture, both of which vary substantially between anatomic sites and across individuals. A promising technique that takes bone microarchitecture into account is microstructural finite element (microFE) analysis. A very large number of finite elements is needed to accurately represent a human bone with its intricate microarchitecture; hence, the resulting microFE models possess a very large number of degrees of freedom.

Detailed microFE models are obtained through high-resolution micro-computed tomography (microCT) of trabecular bone specimens. This allows non-destructive imaging of the trabecular microstructure in living patients with resolutions on the order of 80 microns.

The discrete formulation is based on a standard finite element (voxel) discretization for linear elasticity. The degrees of freedom represent the dis-

placements in the three axis directions at the voxel vertices.

A typical problem size is 100 million voxels corresponding to about 300 million degrees of freedom. Higher resolutions or larger bone specimens lead to

conditioned conjugate gradient algorithm. We employ an algebraic multigrid (AMG) preconditioner based on smoothed aggregation. A particular strength of our parallel finite element solver, ParFE, is the ability to construct the preconditioner without actually




Figure 1: Effective strain in a vertebral human bone specimen. This model of about 1.6 billion degrees of freedom was solved on an IBM BG/L. It is the largest model we solved so far. The visualization was done by the Data Analysis & Visualisation Unit at CSCS, the Swiss Supercomputing Centre.

larger models. We have recently simulated bones with up to 1.6 billion degrees of freedom (see Figure 1). A problem of 1 billion degrees of freedom generates a data volume – essentially the system matrix – of more than 1.25 TB, which must be stored in main memory.

We solve the resulting symmetric positive definite linear systems by the pre-

forming the system matrix. This matrix-free procedure reduces the aforementioned memory requirements by 75%, meaning that only a quarter of the processing elements is needed to solve the problem. Our rule of thumb is half a million degrees of freedom per gigabyte of main memory. We use ParFE on the Cray XT3/XT4 and on the IBM Blue Gene to solve real-world problems,




Figure 2: Repartitioning of a highly porous model of 1.5 million nodes with sixteen processors; (left) the initial linear partition; (right) a well-balanced repartition.

with the typical time to solution being just a few minutes. In contrast to pure bone computations, implants entail variable material parameters. Bone grafts and biomaterials are often used to aid the repair of complicated fractures. The healing process can be analysed using recently developed high-resolution in vivo peripheral quantitative computed tomography (pQCT). While pQCT is able to provide detailed information on bone structure, it cannot do the same for mechanical stability. However, such information can be derived from microFE analyses, and our solver can handle this multi-material situation without problems. Stiffness values as assessed from microFE analyses correlated well with the experimentally measured stiffness.

A crucial issue in large-scale computations is load balancing. This holds true

in particular for complicated structures such as trabecular bones. The layered data layout that is induced by the CT data is not appropriate for the structural analysis, but must be redistributed in order to reduce interprocessor communication (see Figure 2). This portion of the code, that is the repartitioning, does not yet scale beyond 1000 processors. We need to resolve this issue before we can efficiently tackle problem sizes of several billion degrees of freedom. We are also working on an entirely new approach that exploits the underlying Cartesian grid in order to save more memory space.

A further important issue is the visualization of the results. This work is done interactively and hence takes much more wall clock time than the simulation itself. The large data volume again requires parallel processing.

This project is a collaboration between the Institute of Computational Science and the Institute for Biomechanics (both at ETH Zürich), the Swiss Supercomputing Centre (CSCS) and the IBM Zurich Research Center.

Links:

<http://people.inf.ethz.ch/arbenz/projects/bone.html>
<http://parfe.sourceforge.net>

Please contact:

Peter Arbenz
 Institute of Computational Science,
 ETH Zürich
 Tel: +41 44 632 74 32
 E-mail: arbenz@inf.ethz.ch

Ralph Müller
 Institut für Biomechanik, ETH Zürich
 Tel: +41 44 632 45 92
 E-mail: ram@ethz.ch

Using Desktop Grids to Securely Store e-Health Data

by Jesus Luna, Manolis Marazakis and Marios D. Dikaiakos

A collaboration between CoreGRID partners FORTH (Foundation for Research and Technology – Hellas) and UCY (University of Cyprus) is investigating the design of Secure Desktop Grids that are compliant with the strict EU data protection legislation set for e-Health applications. The objective of this effort is to explore the use of a desktop Grid to store the data generated by the Intensive Care Grid System (ICGrid), developed by UCY and the Nicosia General Hospital (Intensive Care Unit).

Desktop Grids are a specialized form of distributed system, whereby shared resources (processor or storage) are provided on a voluntary basis by resource

contributors. These environments are able to provide commodity resources not only for CPU-intensive tasks, but also for applications that require signif-

icant amounts of memory, disk space and network throughput. For example, the widely used Berkeley Open Infrastructure for Network Computing

Annotation Parameters

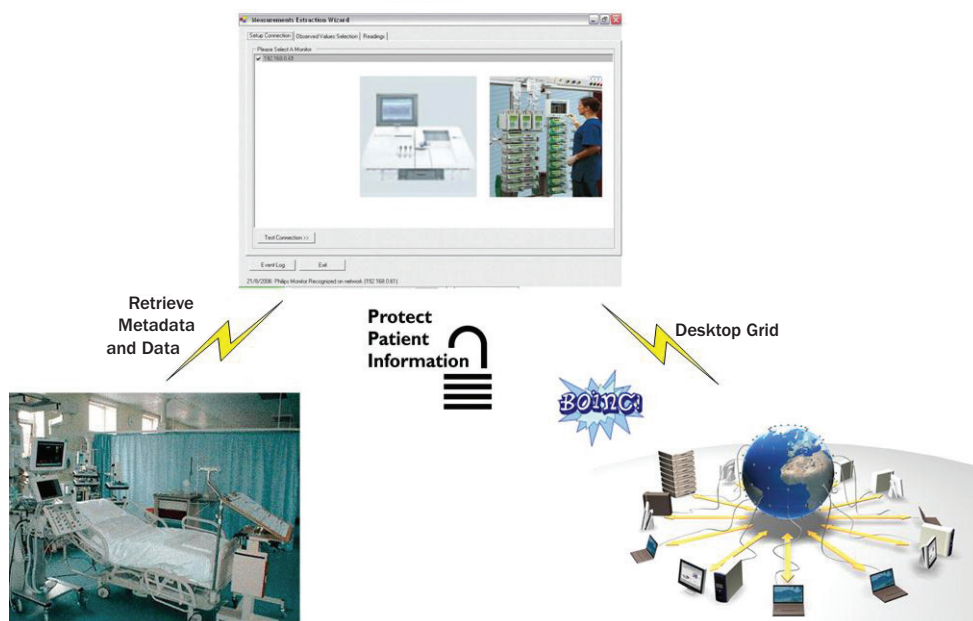


Figure 1: While the storage capability of desktop Grids is quite adequate for the requirements of e-Health systems, critical security and privacy issues must first be addressed.

(BOINC) infrastructure can provide a sustained rate of 95.5 teraflops, 7.74 petabytes of storage and an access rate of 5.27 terabytes per second.

The potential computing power of desktop Grids is, however, even greater. With the number of Internet-connected PCs projected to reach 1 billion by 2015, a distributed computing power will be in place with several petaflops and a storage capacity of around one exabyte, far exceeding what can be provided by any centralized system. It is therefore unsurprising that a significant amount of effort is now aimed at 'volunteer computing' as a new paradigm for both the computational and data Grids.

On the other hand, modern e-Health systems require advanced computing and storage capabilities, leading to the adoption of technology like the Grid and giving birth to health Grid systems. In particular, intensive care medicine uses this paradigm when facing the high flow of data coming from intensive care units (ICUs). Sensors monitoring an ICU inpatient typically generate around 700MB of data per day, in addition to images, physician's annotations and so forth. These data must be stored so that, for example, data-mining techniques can be applied at a later time to find useful correlations for practitioners facing similar problems.

Unfortunately, moving an ICU patient's data to a desktop Grid requires an inte-

gral security solution to be established. This must be harmonized with the current EU data protection legislation and be able to avoid common attacks on the data and metadata being managed. There is a clear need not only to identify the vulnerabilities associated with health Grids and desktop Grids, but also to design new mechanisms able to provide confidentiality, availability and integrity. To this end, the joint research taken early this year by the CoreGRID's WP2 partners, FORTH (Greece) and UCY (Cyprus), has proved that the greatest threat to patient privacy comes in fact from desktop clients, which cannot be fully trusted and may easily leak personal data. In an effort to cope with the identified privacy gaps, this joint research has also contributed a security protocol for providing privacy to health Grid systems from a data-centric point of view, while maintaining compatibility with EU legal frameworks and high-level security mechanisms (ie electronic health cards). The proposed privacy protocol makes use of two well-known data-security techniques (encryption and fragmentation), plus a novel concept involving a quantitative evaluation of the security level associated with the desktop Grid's nodes.

We expect to implement this novel protocol in the coming months within the ICGrid system, using the gLite middleware and paying particular attention to performance issues related to the use of cryptographic mechanisms.

Links:

<http://grid.ucy.ac.cy>
<http://www.ics.forth.gr/carv/scalable>
<http://boinc.berkeley.edu>
<http://www.edges-grid.eu>
<http://www.coregrid.net>

Please contact:

Jesus Luna
 University of Cyprus
 Tel: +357 22892663
 E-mail: jluna@cs.ucy.ac.cy

Manolis Marazakis
 FORTH-ICS, Greece
 Tel: +30 2810391699
 E-mail: maraz@ics.forth.gr

Marios D. Dikaiakos
 University of Cyprus
 Tel: +357 22892700
 E-mail: mdd@cs.ucy.ac.cy

Supercomputing in Clinical Practice

by Stefan Zasada, C.V. Gale, Steven Manos and Peter Coveney

We introduce the GENIUS project, part of the revolution in on-demand access to large-scale computing resources for cerebrovascular modelling, surgical planning and intervention.

Traditional medical practice requires physicians to use their judgement and experience to decide on the course of treatment best suited to an individual patient's condition. While training and practice help physicians to hone the ability to choose the most effective treatment for a particular ailment, this decision-making process often fails to take into account all of the data potentially available. Indeed in many cases, the sheer volume or nature of the data available is too much for a human to analyse, and valuable information is discarded from the decision-making process.

The GENIUS project (Grid-Enabled Neurosurgical Imaging Using Simulation) is funded by the UK EPSRC and coordinated by University College London, and involves the Centre for Computational Science, University College London and neurosurgeons at the National Hospital for Neurology and Neurosurgery (Queen's Square, London). The aim of the project is to augment the clinical decision-making process through the development of patient-specific models based on image data, to simulate blood flow within a given patient's brain. These simulations are used in the planning of surgical intervention for arteriovenous malformations and aneurysms. Crucially, this project makes use of large-scale computing resources through on-demand access – a key requirement if patient-specific simulation is to be of use in clinical practice.

Briefly, this project aims to achieve the following scenario. A three-dimensional rotational angiography scan of a patient's brain vasculature is taken in London; a blood-flow simulation is then performed with resources made available on the US TeraGrid, or in the UK (HeCTOR, HPCx or NGS); results in the form of blood-flow visualization are returned in real time to a workstation in the operating theatre control room (see Figure 1) and are either analysed on the fly for urgent surgical intervention, or later as required, for

prognostic and/or surgical planning purposes (see Figure 2).

In practical terms, the patient-specific medical simulation scenario described above requires access to both appropriate patient data and to the infrastructure on which to perform potentially very large numbers of complex and demanding simulations. Resource providers must furnish access to a wide range of resources, typically made available through a computational Grid, and must institute policies that enable patient-specific simulations to be run on those resources. Unfortunately, the traditional 'batch queue' model used by the majority of high-performance computing (HPC) resource providers does not meet the medical computing requirements of this project and others like it. Simulations used to support clinical procedures must be run when required by the clinician: surgical procedures cannot be postponed until a simulation reaches the top of a batch queue. Emergency medical simulations must therefore either pre-empt all other jobs running on a machine, or be scheduled into the clinical workflow through advance reservation.




Figure 1: Cerebrovascular blood flow – von Mises stress field. Blue indicates a stress of 0 Pa and red a stress of 6.7 Pa.

Consequently, the GENIUS project is trialling technology such as HARC and SPRUCE (see links below) that allows batch-processing HPC resources to be used in a more interactive manner. HARC allows users to co-reserve time across multiple resources, which is essential not only when multi-site Message Passing Interface (MPI) jobs need to be run (eg using MPIg; see links below), but also to the scheduling of simulations into the clinical timetable. SPRUCE in contrast allows a user to pre-empt jobs that are already running or are queued on a machine, when it becomes necessary to run an 'emergency' job. Traditionally, such 'urgent' access to computational resources has only been made available in the US TeraGrid for weather-related simulations (eg in the event of a hurricane).

Apart from the requirement for resource providers to support interactive and emergency jobs on their machines, the ability to routinely run computational simulations in support of medical treatment procedures will also require those funding bodies involved in supporting medical research to engage with the HPC community, and develop policies on the provision of HPC access to underpin this growing field of research. Key issues include the impact the requirements of such an infrastructure have on the policies put in place by Grid resource providers, and why resource providers must listen to the requirements of their current user communities and formulate policies that meet the needs of those communities, while also tracking the needs of emerging user groups.

Computational biomedicine, the field within which the GENIUS project operates, is gaining momentum through projects such as the EU FP6 e-Health ViroLab project and the FP7 Virtual Physiological Human Initiative. The latter receives EU funding of €72 million and involves numerous universities, hospitals, organizations and industrial partners throughout Europe. To date however, efforts in this realm have been hampered by a lack of 'joined-up'

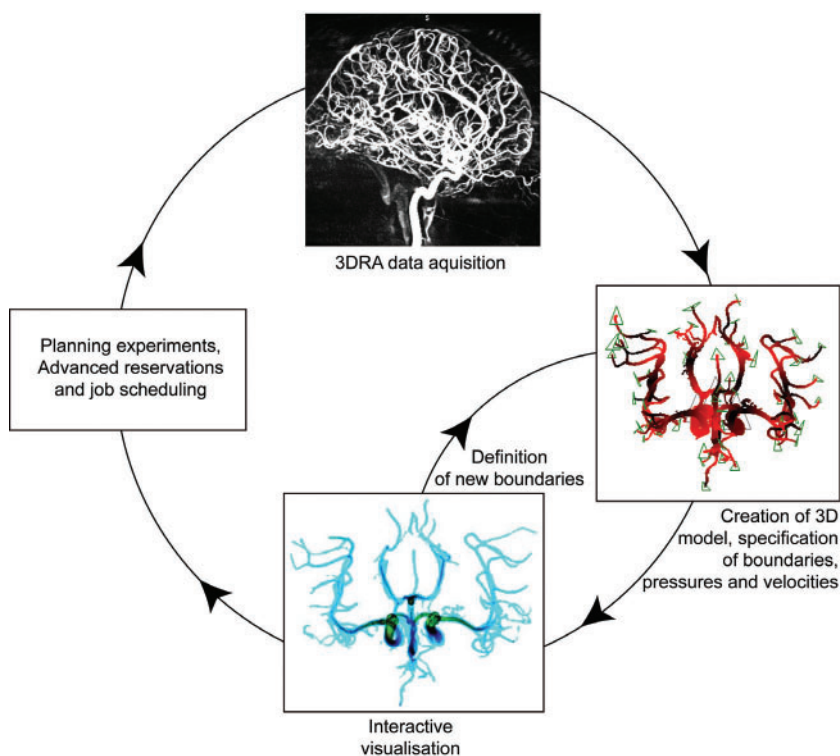


Figure 2: Workflow diagram of GENIUS clinical 'scenario'.

resource provision strategies, at national and international levels, to allocate appropriately scaled resources to projects when they are funded. Such projects often require access to a wide range

of resources as part of a scientific workflow, for example the high-end machines provided by DEISA (Distributed European Infrastructure for Supercomputing Applications) or the TeraGrid, along

with smaller clusters provided by the UK National Grid Service.

The GENIUS project will be complete in December 2009, at which time a unique test case for the use of patient-specific simulation and high-performance computing resources to plan surgical intervention will be available. As described, however, many challenges remain to be resolved before supercomputing at work becomes the premise of the medical doctor.

Links:

- GENIUS project Wiki: <http://wiki.realitygrid.org/wiki/GENIUS>
- VPH Initiative and VPH NoE: http://www.biomedtown.org/biomed_town/vphn
- MPICH-G2: <http://www3.niu.edu/mpi/>
- SPRUCE: <http://spruce.teragrid.org>
- Urgent Computing, CTWatch: <http://www.ctwatch.org/quarterly/archives/march-2008>

Please contact:

Peter Coveney
 Centre for Computational Science,
 Department of Chemistry, University
 College London, UK
 Tel: +44 20 7679 4560
 E-mail: p.v.coveney@ucl.ac.uk

Large-Scale Immune Models and Visualization

by Dimitri Perrin and John Burns

Large-scale computing architectures allow detailed modelling of complex systems. Here we present applications in the field of computational biology.

In silico simulation methods – simulations within computer software – have become indispensable tools in the development of expensive new technology, from aircraft manufacture to nuclear power station development. For various reasons however, modelling and simulation techniques have only very slowly been adopted by the biological research community.

Biological systems are typically complex and adaptive. Given the dynamic nature of these phenomena, it is non-trivial to provide a comprehensive description of such systems and, in particular, to define the importance and

contribution of low-level unsupervised interactions to the overall evolution process. An agent-based approach is presented here.

The immune response, in particular, consists of relatively simple cell-level interactions, but is capable of highly diverse and individual behaviour, with outcomes that are self-evidently more than the sum of their parts. In what follows, a one-to-one reciprocity between autonomous agents and immune cells is proposed. The emergence of tissue-level and body-level patterns is clearly dependent on the relative sizes of cell populations and the balance these

achieve. This is demonstrated, for example, by the course of HIV progression, where the whole immune system collapses once immune cell counts decrease below critical levels.

High temporal granularity is necessary to realistically account for cell mobility and interactions. With current computing resources, this prevents models accounting for every immune cell of every type over the whole body. A compromise between agent diversity and agent population size is therefore required. We have implemented a flexible model, which currently includes viral agents and an agent type for each

entity of the cell-mediated response. This guarantees enough diversity to account for principal immune and viral mechanisms, while permitting the simulation of large populations.

In this way, cell mobility can be examined. Agents are explicitly located in structures modelling lymph nodes. Agent movement within these leads to the occurrence of cell-level interactions, and their passage from one node to another. Thus, lymphatic chain structures allow for the spread of viral infection and connected immune response. The model successfully accounts for rapid spreading throughout the initial lymph chain (ie within hours), and slower progression through the whole body (weeks being necessary to obtain infection of all nodes).

Using parallel implementation on a large cluster computer, simulations can model up to a thousand lymph nodes for a total of more than one billion agents. This permits accurate modelling of the immune response to infections, and inclusion of localized effects such as HIV early infection within the gastrointestinal tract, whose importance was recently highlighted. Equivalent large-scale frameworks are also applicable to other complex systems, biological or otherwise. In particular, *in silico* techniques have enormous potential to contribute to the quality and cost-effectiveness of the drug development process.

Visualization of the results from such simulations is non-trivial. In what follows, we present some initial results in the visualization of immune responses that allows the user to parameterize their *in silico* experiments in order to study the maximum infection clearance dynamics. The tool also proves to be a useful aid in the classroom for students in their early undergraduate training.

Collaboration between ITT Dublin and Dublin City University has led to the development of novel visualization techniques to study cellular interaction from both the large-scale spatial and temporal perspectives. Visualization of processes offers the researcher the ability to speed up, slow down and hypothesize various parameters. Visualization aids understanding as we rely on visual perception to make crucial decisions. For example, with our initial model, we can visualize the dynamics of an ideal-

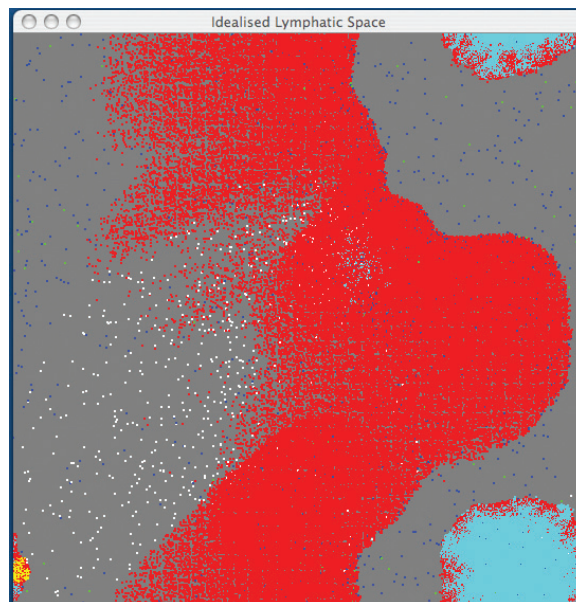


Figure 1:
Idealised lymphatic space.

ized lymphatic compartment, with APC and CD8 cells.

We can parameterize many aspects, such as initial cell levels, rates of change in the life cycle, frequency of infection events and many others. Some viral pathogens are capable of persistent reinfection in that, although population levels of infected antigen presentation cells may decline in response to clearance pressure by a specific CD8 response, over time the number of infected cells rises to chronic and sometimes acute levels. Examples of such viruses are HIV, Human T-lymphotropic virus (HTLV) hepatitis C (HCV), hepatitis B, cytomegalovirus CMV, Epstein-Barr virus (EBV) and rubella.

From the figure we see that as the primary response continues, the effector cells now begin their exit from the lymphatic compartment, after which they will be carried through the blood system and will migrate towards the site of the initial infection (if one exists).

We have developed a 'front end' visualization component to allow students and lecturers in the classroom and lab to experiment with a variety of parameters. We are strongly motivated in this research by the findings of a recent EU report indicating that *in silico* modelling and simulation of biological processes is a key requirement in the development of cost-effective and timely new disease therapies.

Link:

<http://sci-sym.computing.dcu.ie>

Please contact:

Dimitri Perrin
Dublin City University
Tel: +353 1 700 8449
E-mail: dperrin@computing.dcu.ie

John Burns
Institute of Technology Tallaght,
Ireland
Tel: +353 1 404 2766
E-mail: john.burns@ittdublin.ie

High-Performance Computing for Modelling Bacterial Communities

by James T. Murphy, Ray Walshe and Marc Devocelle

Micro-Gen is a tool for modelling the population dynamics of bacterial colonies by taking into account the unique characteristics of each individual bacterial cell and its local environmental conditions. This level of granularity requires significant computing resources when modelling large numbers of agents. However, by implementing an efficient parallel algorithm, Micro-Gen can take advantage of high-performance computing resources to scale up to biologically realistic numbers of bacterial agents while maintaining information about each individual cell.

Micro-Gen is an 'agent-based' model of bacterial growth and interactions in simulated laboratory culture conditions (see ERCIM News, No. 73, p 39-40). It uses information about each individual bacterial cell (the 'agent') to build up a picture of the population dynamics of the colony as a whole. By concentrating on the characteristics of the cells and allowing the behaviour of the colony to emerge as a product of the random interactions between the independent cells, insights can be gained into the mechanisms that lead to, for example, the development and spread of antibiotic resistance in bacteria.

However, a challenge of modelling from the individual cellular level up is that it requires significant computational resources in order to scale up to biologically realistic numbers of cells. For example, concentrations of bacterial cells (per millilitre) in nature can range from approximately one million cells in seawater to greater than a billion cells in the human gut. Despite rapid advances in traditional desktop computers, their memory and performance limitations mean it is currently not feasible to model the numbers of agents typically found in the human gut on a single processor. Therefore, it is important to be able to take advantage of parallel high-performance computing resources in order to tackle these problems.

Micro-Gen was designed to take advantage of parallel computing resources, where multiple computers/nodes work together in order to solve a complex problem. In the case of Micro-Gen, the simulated environment in which the bacteria exist is divided equally among the available processors, each of which is responsible for a specific subsection of the population (see Figure 1). Communication between separate processors is handled by the message-passing inter-

face (MPI), which is a common communications protocol used in high-performance computing. For example, when a bacterial cell reaches a border separating parts of the environment controlled by two different processors/computers, it is sent across to the other computer in a manner analogous to sending a letter between two postal addresses.

Running simulations in tandem across multiple processors/computers overcomes the traditional limitations of the computational resources found in the desktop space. The model can be scaled up according to the number of processors/computers available in order to increase the environment size from the microscopic level up to a visible scale. However, it is important to optimize

parallel algorithms so as to avoid the problem of diminishing returns as the number of processors is increased. This problem occurs because with more processors there is more traffic (or 'letters') between them, which can result in significant bottlenecks.

For this reason, particular emphasis has been placed on optimizing the communication strategy for sending bacterial agents between different computer nodes in Micro-Gen, while ensuring the integrity of the information in transit. Advances in communication technology such as the availability of high-speed interconnects ('Infiniband') between the computer nodes and optimized communication strategies have minimized the overheads associated with communication traffic. This has




Figure 1: Screenshot of Micro-Gen simulation running in parallel on four separate computing nodes, with bacterial colonies (yellow) growing on nutrient agar medium (blue). Lighter shades of blue represent higher nutrient concentrations.

allowed Micro-Gen to be scaled up in testing to 400 processors.

As a consequence, bacterial populations of several billion agents can now be simulated, and important investigations into the population dynamics of large bacterial populations carried out. Questions such as the effect of scale and the interactions between different sub-populations of bacteria can be investigated and correlated with more traditional mathe-

tical modelling approaches that look at the population as a whole. This would allow investigators to trace back from the behaviour of the entire population to that of its individual components. Evolutionary questions could also be asked by incorporating genetic components into the bacterial agents and simulating the competition between genetically differing strains. For further information about the biological aspects of Micro-Gen see ERCIM News, No. 73, p 39-40.

Links:

<http://www.computing.dcu.ie/~jamurphy/>
<http://sci-sym.computing.dcu.ie/>

Please contact:

James T. Murphy,
 School of Computing,
 Dublin City University, Ireland.
 Tel: +353 1 700 8449
 E-mail: jamurphy@computing.dcu.ie

Harnessing the Power of Supercomputing

by Richard Blake

The explosive growth in computing power has brought computers into the heart of most scientific disciplines. While servers and workstations service this mainstream, it is parallel supercomputers capable of harnessing thousands of processors for a single calculation that are redefining the boundaries of computational science. The Science and Technology Facilities Council's (STFC) Computational Science and Engineering Department (CSED) is at the forefront of efforts to use supercomputing to tackle problems across a range of scientific disciplines.

Spread across STFC's Daresbury and Rutherford Laboratory sites, the CSED has a long pedigree of working in high-performance computing (HPC). The heart of our work is the development and application of powerful simulation codes, usually in collaboration with university research groups. Through our involvement in the Collaborative Computational Projects, we bring together leading UK expertise in key fields of computational research to tackle large-scale scientific software development. Our work spans a range of disciplines from quantum chemistry through to biology and engineering, allowing the latest techniques in HPC to be brought to bear in many different areas.

As well as developing our own codes, we also help others redesign their codes to run on the latest machines or take advantage of newly emerging technology such as FPGAs and co-processors. As supercomputing moves into the era of the petascale (systems running at around 10¹⁵ operations per second) we are actively engaged in research to understand the challenges that programming such massive machines will bring. We are currently involved in both of the UK's major academic supercomputers, HPCx and HECToR. The 55 teraflop/s (one teraflop/s equals 10¹² floating point operations per second) HECTOR machine provides the power to run the

largest calculations, and we are looking at innovative ways to use the 12.9 teraflop/s HPCx as a complementary resource to Hector, taking advantage of emerging technology such as Grid computing or using it to drive large-scale visualizations.

To give a flavour of the kind of research that supercomputers enable, we here highlight two projects in which we are involved: one in climate modelling and the other in the development of quantum computing.

Modelling the Oceans

The Proudman Oceanographic Laboratory Coastal Ocean Modelling System (POLCOMS) is a simulation code that has been designed to model the global marine environment. At its core, POLCOMS consists of a 3D hydrodynamic model that simulates the effects of pressure, density etc to provide the forces that drive the flow of water. POLCOMS is an interdisciplinary code, so the core can be integrated with other models such as ecosystem (phytoplankton, zooplankton etc) or wave models. Integrating the models or increasing the resolution of any of the components can increase the computational demands by factors of ten to forty, so extensive work has been carried out on the optimization and parallelization of POLCOMS.

Today POLCOMS can be run on thousands of processors on machines such as HECToR. Work is currently underway to enable POLCOMS to be run concurrently on multiple supercomputers, fully utilizing the total computing power available across the machines. Recent studies carried out with POLCOMS have investigated the northwest European shelf seas. The study has been able to separate the model currents into those driven by winds, by changes in density and by the North Atlantic. The study has demonstrated that, contrary to previous thinking, the long-term circulation is actually driven by the smaller but directional density currents and not the stronger, variable winds.

Designing the Next Generation of Supercomputers

Quantum computing promises a revolution in the capabilities of computers, harnessing the bizarre properties of matter on a quantum scale to enable calculation rates that are inconceivable with the current computer designs.

The materials group within CSED has been investigating the spin interactions in carbon peapods, a novel material that holds promise as a component of a quantum computer.

Carbon peapods are carbon nanotubes containing tiny spherical balls of car-




Figure 2: Charge transfer in a Sc@C82@(14,7) peapod as calculated by DFT with CRYSTAL.

bon known as fullerenes. Inserting a metal atom such as scandium into each fullerene introduces an unpaired electron that turns each fullerene into a tiny magnet capable of interacting with its neighbour. Each magnet can hold a qubit (the fundamental unit of quantum computing), so a chain of these magnets interacting together is a potential candidate for the core of a quantum computer.

The calculations to model the spin interactions require the use of the computationally demanding density functional theory (DFT). As the carbon

peapod system contains more than 600 atoms, this sort of calculation is intractable with standard codes on a conventional system.

The group has therefore been using the CRYSTAL code developed jointly by CSED and the University of Turin to run the calculations on HPCx. It is only by using the power of such large machines with codes like Crystal that these sorts of calculations become possible.

Links:

- CSED: <http://www.cse.scitech.ac.uk>
- HECTOR: <http://www.hector.ac.uk>
- HPCx: <http://www.hpcx.ac.uk>
- CRYSTAL: <http://www.cse.scitech.ac.uk/cm/g/>
- CRYSTAL POLCOMS: <http://www.pol.ac.uk/home/research/polcoms>

Please contact:

Richard Blake (Acting Director)
 Computational Science and
 Engineering Department
 Daresbury Laboratory, UK
 E-mail: r.j.blake@dl.ac.uk




Figure 1: Summer mean horizontal circulation around the UK averaged at mid-depth (20-40m) as calculated with POLCOMS.

Fluids and Supercomputers: The Billion Particle Era

by Petros Koumoutsakos

A century of advances in numerical methods, integrated with advances in software and hardware, provide us today with unprecedented tools to study and control flow as it pertains to key problems of our society. This includes energy (wind turbines, aircraft wakes), health (microfluidics, hemodynamics) and nanotechnology (nanofluidics, nanomedicine).

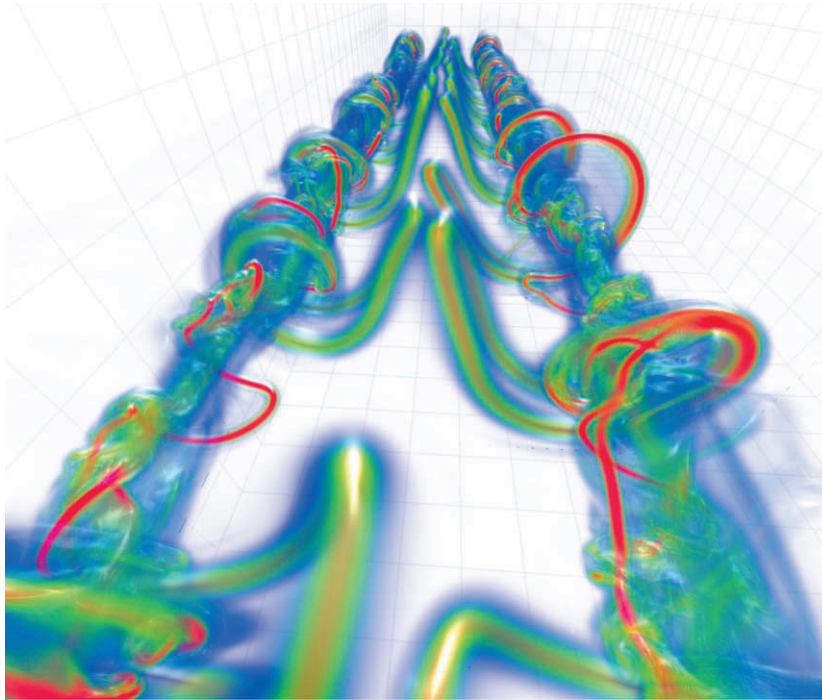


Figure 1: Visualisation of aircraft vortex wakes undergoing a long wavelength instability.

At the Chair of Computational Science at ETH Zurich, we are developing multiscale particle methods for simulating diverse physical systems [1], with an emphasis on their implementation in high-performance computing (HPC) architectures. Multiscale simulations often require a level of algorithmic complexity that is difficult to translate into effective parallel algorithms. A key aspect of our approach is the development of numerical methods with an emphasis not only on their accuracy and stability but also on their software engineering and implementation on massively parallel computer architectures.

In one such example from the last two years, we have developed particle methods for the study of vortex-dominated flows. We collaborated with researchers from the IBM Zurich Research Laboratory to port these methods onto massively parallel computer architectures. This enabled state-of-the-art direct numerical simulations of aircraft vortex

wakes [2], which employed billions of vortex particles and exhibited excellent scalability for up to 16,000 nodes of the IBM Blue Gene/L (BG/L).

Resurrecting Vortex Methods for HPC

This effort represents the culmination of fifteen years of research into vortex methods [2], and has combined a number of ingredients that only a few years ago were considered unsuitable for flow simulations.

We first resurrected a numerical method (vortex methods) with Lagrangian adaptivity and a minimal use of computational elements for vortical flows. Vortex methods were in fact the first computational fluid dynamic (CFD) technique employed to study fluids in the 1930s (vortex sheets were computed by hand by Rosenhead and his human 'computers'), and were also used in the first digital computer flow simulations in the 60s and 70s. Vortex methods have since fallen out of favour due to their inaccuracy, which is largely attributed

to the distortion of the Lagrangian computational elements.

We then addressed the low memory available in massively parallel computer architectures such as the BG/L. This architecture was considered unsuitable for flow simulations requiring large per-processor memory allocations and extensive communications when solving for elliptic differential equations, as imposed here by our velocity-vorticity formulation.

The solution to this multiobjective problem was found by reformulating several aspects of vortex methods: coupling particles with grids to regularize their locations, developing novel data structures for the distribution of the particles, and adopting fast Poisson solvers effectively implemented in the BG/L architecture. Last but not least, the development of software libraries and open-source software (see our lab Web site) has enabled the sustained development of these tools over a number of years.

How do you look at two billion particles?

Massively parallel simulations using up to two billion particles have provided us with a unique insight into the development of long-wavelength instabilities in aircraft wakes.

At the same time, a new challenge has emerged, namely, the visualization and processing of the large data sets that become available through the simulations. Each snapshot of the flow field results in about 100 Gbytes of data, meaning the usual techniques such as flow animation and walkthrough visualization in turn become computationally intensive tasks. The continuum flow field needs to be reconstructed from the particle properties, and high-performance visualization techniques are necessary in order to process the data. We are pursuing this approach in collaboration with scientists at the Swiss National Supercomputing Center (CSCS) in Manno, aiming to provide a unique look at the intricacies of vortex structures in aircraft wakes. A parallel approach involves the use of multiresolution techniques such as wavelets, leading to in situ analysis and structure-informed visualization.

From Insight to Optimization

The increased capabilities offered by high-performance computing enable us to think beyond simulating flows to actually controlling and optimizing them. In our latest collaboration with IBM Zurich and CSCS, we employ evolutionary algorithms in order to discover vortex wake configurations that lead to fast instability growth and wake destruction. Evolutionary algorithms are well suited to parallel computer architectures, but at the same time require a large number of iterations, an issue that makes them seemingly unsuitable for expensive CFD calculations (contemplate thousands of simulations involving billions of computational elements!). We have addressed this problem by exploiting the inherent parallelism of these methods and employing machine-learning techniques to increase their convergence. We are currently discovering vortex configurations that can assist engineering intuition in designing next-generation aircraft.

Next Steps

Our particle framework is being extended to problems ranging from geomechanics to cancer treatments and

nanofluidics. We believe in the integration of computer science expertise with mathematics and application domains, and in collaborative platforms that will become essential for effective high-performance computing. We must be reminded that "τα παντα ρει" (everything flows/ changes) even in the case of supercomputing but we believe that revisiting the sciences with the HPC lens will provide us with unexpected treasures.

Links:

<http://www.cse-lab.ethz.ch>

[1] Multiscale Flow Simulations Using Particle Methods:

<http://arjournals.annualreviews.org/doi/abs/10.1146/annurev.fluid.37.061903.175753>

[2] Billion Vortex Particle Direct Numerical Simulations of Aircraft Wakes: <http://dx.doi.org/10.1016/j.cma.2007.11.016>

Please contact:

Petros Koumoutsakos

Chair of Computational Science, ETH Zurich, Switzerland

Tel: +41 44 632 52 58

E-mail: petros@ethz.ch

Supersonic Flow Simulation on Emulated Digital Cellular Neural Networks

by Sándor Kocsárdi, Zoltán Nagy, Árpád Csík and Péter Szolgay

In the areas of mechanical, aerospace, chemical and civil engineering the solution of partial differential equations (PDEs) has long been a primary mathematical problem. In this field, one of the most exciting areas of development is the simulation of fluid flow, which involves for example problems of air, sea and land vehicle motion. The governing equations are derived from the Navier-Stokes equations and solved by using first and second-order numerical methods. In collaboration with the Department of Mathematics and Computational Sciences, Széchenyi István University and the Department of Image Processing and Neurocomputing at the University of Pannonia, researchers at the Cellular Sensory and Wave Computing Laboratory of SZTAKI are working on finding an optimal computational architecture that satisfies the functional requirements with the minimal required precision, while driving computing power toward its maximum level.

A wide range of industrial processes and scientific phenomena involve gas or fluids flowing over complex obstacles, eg the flow of air around vehicles or buildings, or that of water in the oceans. In engineering applications the temporal evolution of non-ideal, compressible fluids is often modelled by the system of Navier-Stokes equations. By neglecting

all non-ideal processes and assuming adiabatic variations, we obtain the Euler equations, describing the dynamics of dissipation-free, inviscid, compressible fluids. They are a coupled set of nonlinear hyperbolic partial differential equations and form a relatively simple yet efficient model of compressible fluid dynamics. Unfortunately, the

necessity of the coupled multi-layered computational structure with nonlinear, space-variant templates does not make it possible to utilize the huge computing power of the analogue Cellular Neural Network Universal Machine (CNN-UM) chips. To improve the performance of our solution, emulated digital CNN-UM implemented on field pro-

programmable gate arrays (FPGAs) has been used. Thus, we intend to perform the operations with the highest possible parallelism.

Since the logically structured arrangement of data is fundamental for the efficient operation of the FPGA-based implementations, we consider explicit second-order accurate finite-volume discretization of the governing equations over structured mesh employing a simple numerical flux function. The main advantage of this method over the forward Euler method, which is used extensively in the computation of the CNN dynamics, is that this approximation is more robust in the case of complex computational geometries and in the presence of shock waves in the solutions. Indeed, the corresponding rectangular arrangement of information and the choice of a multi-level temporal integration strategy ensure the continuous flow of data through the CNN-UM architecture.

The Falcon architecture is an emulated digital implementation of CNN-UM array processor that uses the full signal range model. On this architecture the flexibility of simulators is combined with the computational power of analogue devices. Not only can the template size and computational precision be configured, but space-variant and nonlinear templates can also be used. In accordance with the discretized governing equations, we have designed a complex circuit that is able to update the values of the conservative state vector of a cell in every clock cycle using the emulated digital CNN-UM architecture.

Implementation and testing of such an application-specific arithmetic unit can be very time-consuming. However, using rapid prototyping techniques and high-level hardware description languages such as Handel-C from Agility makes it possible to develop the optimized arithmetic unit much faster than using an approach based on conventional hardware description language (HDL).

To show the efficiency of our solution, we used a complex test case in which the flow profile over a forward-facing step was computed. The simulated region is a two-dimensional cut of a pipe that is closed at the upper and lower boundaries and open at the left




Figure 1: The density values of the Mach 3 flow after 0.4s, 1s and 4s computed in a model having an array size of 320×960.

and right boundaries. The direction of flow is from left to right and the flow speed at the left boundary is constant and set at three times the speed of sound (Mach 3). The solution contains shock waves reflecting from the closed boundaries. Figure 1 shows the results of the computation using the derived method after 0.4s, 1s and 4s of simulation time with a 0.39ms (1/2560s) time step.

The proposed circuit can be implemented on mid-sized gate arrays on Agility RC203 and RC2000 boards. In the case of second-order approximation, an approximately 21-fold speedup can be achieved compared to the Intel Core2Duo T7200 processor running at 2 GHz. By using larger FPGAs, the achievable performance can be improved more than 589-fold.

In future the designed arithmetic unit will be extended to three-dimensional flow problems and non-uniform computational Grids could be possible.

Links:

MTA SZTAKI, Cellular Sensory and Wave Computing Laboratory:
<http://lab.analogic.sztaki.hu/>
 Agility Design Solutions:
<http://www.agilityds.com/default.aspx>

Please contact:

Sándor Kocsárdi
 University of Pannonia, Hungary
 E-mail: skocso@vision.vein.hu

High-Performance Computing of Multiphase Flow

by Bipin Kumar, Yan Delaure and Martin Crane

Appropriate understanding of the behaviour of multiphase environmental flows requires CPU-intensive computational modelling. Large systems of algebraic equations must be solved iteratively and this has motivated the development and implementation of parallel iterative algorithms.

Multiphase flow problems are found in ocean and naval applications, chemical and other industrial and environmental processes, and in a variety of biomedical applications. In ocean engineering, wave-induced forces have a significant impact on structural and general hydrodynamic design. In particular, breaking waves may pose serious safety problems and affect the offshore structural integrity. Computational modelling of wave-induced hydrodynamics is often used in the design of conventional offshore structures in order to evaluate environmental loads. In the case of chemical applications, bubble formation or injection can be used to enhance mass transfer to improve chemical reactions. Bubble formation due to cavitation is also found in a range of pumping applications and may induce structural damage or loss of efficiency. Computational modelling is used to predict these types of multiphase flow phenomena.

The examples mentioned here are of the flow of two fluids separated by sharp interfaces. Numerically modelling

advection of this interface is a challenging task and a range of methods have been developed over the last thirty years. The Volume of Fluid (VOF) method is a widely used approach that approximates the position, slope and curvature of the interface between two fluids using the volume fraction of each fluid. It is this volume fraction that is advected with the flowing fluid. The main difficulty of the VOF method is, to preserve the sharp interface during advection of fluid at the cell face.

Current research in our lab attempts to improve the models for this problem by modifying the VOF method. Although the VOF method is not necessarily CPU intensive, it imposes a further restriction on the solution of the fluid flow equations. The flow is almost exclusively unsteady and the time step generally needs to be significantly reduced by comparison with conventional fluid flow problems. This is particularly the case when surface tension must be taken into account. The sharp changes in the fluid properties at the interface also lead

to algebraic systems that are more challenging to solve.

Mathematically the fluid properties and other phenomena are represented by a set of partial differential equations (PDEs) called the Navier-Stokes (NS) equations. In the process of numerical solution, PDEs are first converted to a large system of linear algebraic equations. The algebraic system of 'n' equations is represented by $Ax = B$, where 'A' is an $n \times n$ matrix and 'x' and 'B' are the n-dimensional vectors.

It is known that there exist two types of solvers – direct methods and iterative methods – for solving algebraic equations. Iterative methods are commonly employed to solve such large systems of equations, and are further classified into stationary and non-stationary types. In the non-stationary category, current variable values are updated based on several previous iteration values. The research being conducted at our lab employs the Krylov Subspace (KS) non-stationary methods. This work will

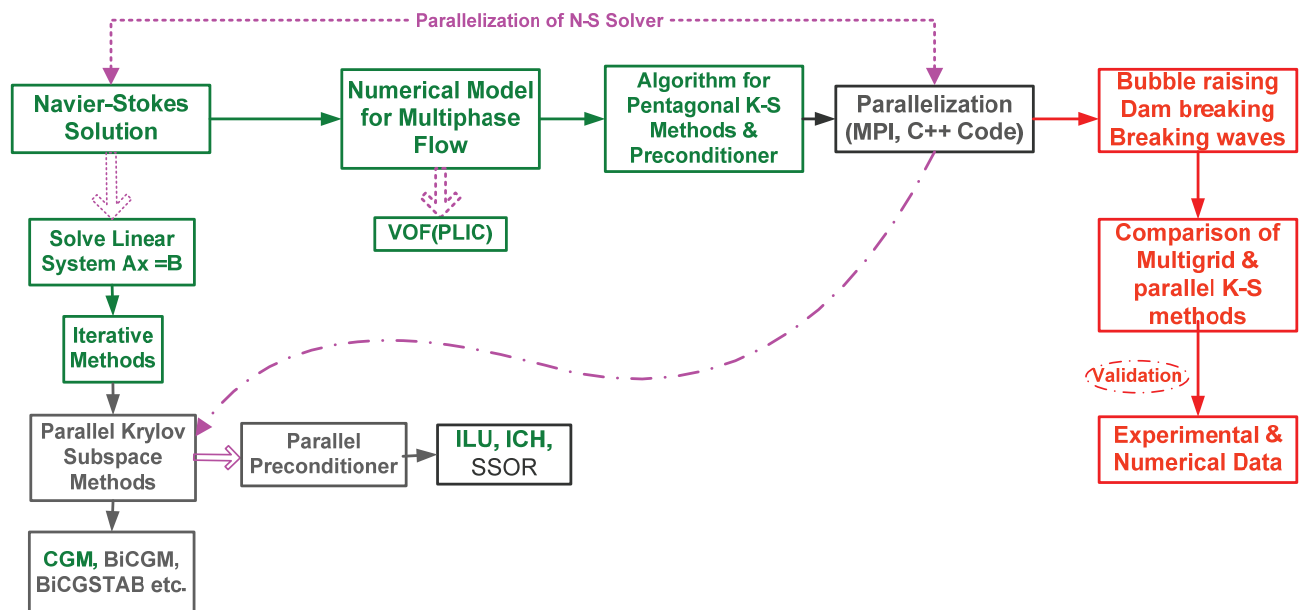


Figure 1: Steps of the multiphase flow modelled on a parallel computer.

investigate the power of some of KS methods such as the conjugate gradient method (CGM), bi-conjugate gradient method (bi-CGM), conjugate gradient stabilized method (CGSTAB) and bi-CGSTAB.

Solving these very large-scale systems is computationally very demanding and time-consuming, and requires a parallel or distributed environment such as a High-Performance Computing Cluster (HPCC).

Due to the sharp change at the interface, matrices arising from PDEs are generally unsymmetric and ill-conditioned. It has been reported that the convergence of KS methods is slow when applied to those matrices. To overcome this difficulty, preconditioning techniques are

applied to matrices in the solution procedure. In this project, preconditioned versions of KS methods have been developed in C++. Further, the parallel version of KS methods and preconditioners have also been developed using a message-passing library (MPI) and C++.

The objective of the project is to develop a numerical model based on the solution of the two- or three-dimensional NS equations to model two-fluid flow. The research focuses both on modelling moving boundaries and on improving the computational efficiency of the iterative methods used to solve the system of algebraic equations. The solution of the NS equations is coupled with the VOF method to simulate the motion of fluid. The parallel version of

KS methods and preconditioners will be implemented on the parallel computer available at Dublin City University, and we propose a series of comparative stability, efficiency and scalability analyses of the parallel solver versus its serial version. The entire simulation process is illustrated in Figure 1.

Links:

http://www.dcu.ie/mechanical_engineering/efmg/projects.shtml
<http://www.computing.dcu.ie/~mcrane/>

Please contact:

Bipin Kumar
 Dublin City University, Ireland
 Tel: +353 17007756
 E-mail: bkumar@computing.dcu.ie

Set-Top Supercomputing: Scalable Software for Scientific Simulations on Game Consoles

by Dimitrios S. Nikolopoulos

Exponential improvements in the performance of computer systems have brought supercomputing to the desks of users. A modern processor has the equivalent computing capacity of a massively parallel computer system from the last decade. The history of Top500, the list of the 500 most powerful computing systems on the planet, indicates that our 2008 laptops would have ranked in the Top500 most powerful computer systems ten to fifteen years ago.

Computer systems architecture has been a catalyst in this amazing development. Until recently, catalysing innovations in hardware would improve software with little or no involvement from users. Even before it hit the market, a new processor would simply run faster software than had previously been written. Unfortunately, this 'free ride' came to an end due to technological barriers. While the complexity of hardware quickly increased to achieve further performance improvement, power consumption also began to increase at an exponential rate, rendering computing systems less energy-efficient and less reliable. Furthermore, exponentially increasing hardware complexity ceased to achieve exponentially improving performance, showing diminishing returns instead.

To address these challenges, the industry turned to processors with multiple cores (known as multi-core or many-core) and off-loaded most of the responsibility for sustaining performance improvement to

software. To exploit multiple cores, software needs to be parallelized, a task which has traditionally been considered hard, time-consuming, error-prone and feasible only for specialized scientific application domains.

The emergence of multi-core processors as a de facto standard brings us to the following interesting situation. A processor sitting in a Sony PlayStation 3, a set-top box that we can buy for around 400 Euros, has nine cores and can perform as many as 200 billion operations per second. Ten years ago, this level of performance would have been achieved only by the fastest computers in the world. To make use of this however, a user needs to invest tremendous effort in porting, parallelizing and specializing code for the particular processor of the PlayStation, the Cell Broadband Engine (Cell/BE).

The real-life experience of the author suggests that it takes over three months


for an advanced PhD student with ample background in parallel computing to port a reasonably sized (a few Klines) computational biology algorithm on this platform. The effort, and perhaps the associated code development methodology, are most likely not portable to other (single-core or multi-core) processors. It is even questionable whether the effort is portable across algorithms and applications running on the same processor.

Researchers at the Computer Architecture and VLSI Systems Laboratory (CARV) of FORTH-ICS are tackling the software productivity crisis that has emerged since the introduction of multi-core processors in the market. A significant component of this research is the development of programming environments (runtime systems, languages, compilers and hardware/software interfaces) and tools for processors with multiple heterogeneous cores. Some of these have a general-purpose

architecture (eg superscalar pipelined), while others have a specialized architecture (eg vector processing units) designed to accelerate specific algorithms and application kernels. Examples of these processors, besides the Cell/BE, are the AMD Fusion processor and the Intel EXO architecture. CARV is also heavily involved in developing next-generation multi-core processors through its participation in the SARC project (Scalable computer ARChitecture).

Our research on software environments for programming current and future heterogeneous multi-core processors focuses on providing a unified framework for exploiting the multiple layers and forms of parallelism available in these processors. We aim at achieving this goal while asking the programmer to use a single 'expression' of parallelism in the program. In object-oriented terms, we are seeking mechanisms for exploiting polymorphism in the expression of parallelism by having programmers use a single notion of algorithmic 'work unit' and map this notion to multiple (overlapping or non-overlapping) physical implementations of parallel execution units. We have developed several research prototypes to support such a 'hardware-independent' framework for parallel programming. These prototypes include:

- An event-driven task execution and scheduling model which achieves better coordinated scheduling of heterogeneous cores, improves the capabilities of software in exploiting simultaneously fine-grained and coarse-grained parallelism and controls dynamically and transparently the execution cores allocated to program components.
- Models of multi-grained (layered) parallelism for systems with many heterogeneous multi-core processors each with many heterogeneous cores. These models drastically prune the design space for mapping software to multi-core computer systems, which is exponential and exhibits non-linear performance effects, when concurrency changes at different layers of the system.
- A runtime program auto-tuning tool, for online empirical searching of program decompositions and mappings to parallel architectures, to optimize performance based on dynamic execution feedback.



```

AAB24882      TYHMCQFHCRCYVNNHSGEKLYECNERSKAFSCPSHLQCHKRRQIGEKTHEHNQCGKAFPT 60
AAB24881      -----YECNQCGKAFQHSSLKCHYRTHIGKPYECNQCGKAFSK 40
                ****: .***: * *:*** * :****: * *****. .

AAB24882      PSHLQYHERHTHTGKPYECHQCGQAFKRCSSLQRHKRHTHTGKPYE-CNQCGKAFQA- 116
AAB24881      HSHLQCHKRHTHTGKPYECNQCGKAFSQHGLLQRHKRHTHTGKPYMNVINMVKPLHNS 98
                ****: *****:****:***: . *****:***** : *.: :

```




Figure 1: Accelerating computational biology algorithms on a PlayStation 3: phylogenetic tree (tree of life, top); sequence alignment (middle); and speedups (bottom) achieved from running three algorithms for phylogenetic inference (RAxML, PBPI) and sequence alignment (SWat) on a PlayStation 3 with one Cell/BE processor using its six accelerator cores, versus an Intel CoreDuo using two Intel Xeon cores and running at 2.8 GHz on an Apple MacBook.

We have deployed the aforementioned components on Sony PlayStation 3 and several other platforms to model and accelerate with little effort several computational biology algorithms, including phylogenetic tree inference and multiple sequence alignment. Figure 1 shows some experimental results, indicating that using the same code basis and our acceleration tools, we are able to use a PlayStation 3 and achieve almost an order of magnitude higher performance than a high-end 2.8 GHz Intel CoreDuo processor on a Mac-

Book. The PlayStation's performance per Euro can be up to two orders of magnitude better than that of the MacBook.

Links:

<http://www.top500.org>
<http://www.sarc-ip.org>

Please contact:

Dimitrios S. Nikolopoulos
 FORTH-ICS, Greece
 E-mail: dsn@ics.forth.gr

Enhancing Traffic Safety by Integrating Real-Time Infrastructure and Vehicle Data in a Cooperative System

by Kashif Din

An Austrian national project carried out by Arsenal Research and MagnaSteyr Fahrzeugtechnik has successfully demonstrated the operation of a cooperative vehicle infrastructure system under real-world conditions. The project was installed on a section of national freeway, and has bidirectional communication facilities, infrastructure-based traffic, weather and road surface data and real-time vehicle (extended floating car) data.

Intelligent transportation systems of the future will take advantage of rapid developments in information and communication technology. Their aim is to make the movement of traffic safer, more efficient, more comfortable and cleaner by linking relevant traffic information coming from the vehicles (eg navigation, in-vehicle data systems) with that from roadside sensor systems (eg traffic flow, density, weather etc). So-called 'cooperative vehicle infrastructure systems' should promote the safe and optimal use of road network capacity and a better response to incidents and hazards, through increased and up-to-date information about vehicles, their location, and road conditions in the whole road network.

The cooperative vehicle infrastructure system used in this project had a total length of six kilometres and was

installed on part of the Austrian freeway A2. The test track was equipped with several wireless access points, enabling a seamless communication zone for bidirectional vehicle-to-infrastructure communication.

The collection of traffic data was performed with a camera system consisting of sixteen video cameras along eight cross-sections. From the camera images, microscopic traffic data such as velocity, type of vehicle (passenger car/HGV) and appropriate lane are derived by image-processing software. Out of this, macroscopic values such as aggregated traffic flow and average velocities for each section and lane are calculated.

Weather data was collected by a weather station installed on the centre of the test track. The station was

equipped with both standard sensors (eg temperature, precipitation type and intensity, wind speed, humidity etc) and additional sensors for precise safety evaluation (eg integrated water film thickness, visibility). Both traffic and weather information were collected through a time-triggered central server system and were prepared for further data processing there.

On the infrastructure side, a complete road surface and road bed measurement was performed with a specially equipped measurement vehicle (Road-STAR), which measured safety-related parameters like skid resistance, transverse evenness, longitudinal profile, roughness, texture and road geometry in a standardized manner.

In order for bidirectional cooperative vehicle-infrastructure interaction to

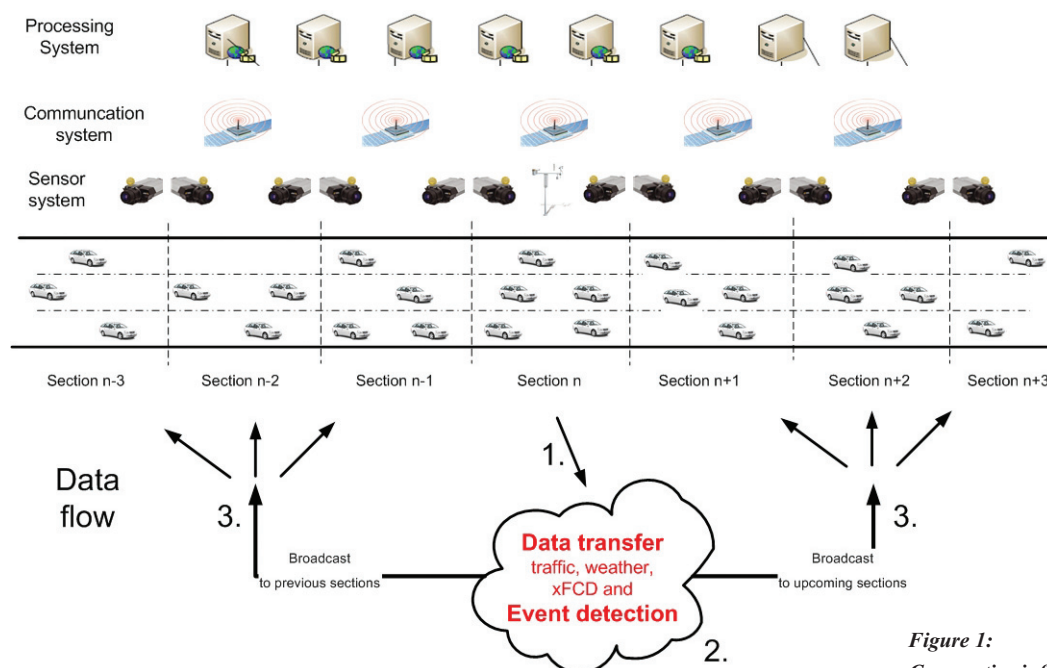


Figure 1:
Cooperative infrastructure vehicle track.

work, it is necessary both to gather information from vehicles driving on the section and to transfer relevant events from them. The basic idea comes from the classical approach of floating car data, where vehicles regularly transmit their positions, which can then be used to generate traffic information (travel times, velocities, level of service).

Within the project, an in-car application was developed that processed vehicle data and detected the following events from it: accident, emergency braking, loss of traction, rain, wet surface and warning light activation. Together with the GPS position, the event code was transmitted to the central system in real time. For demonstration purposes, substitute signals like headlight activation or other 'harmless' events were used to simulate the events mentioned above.

The data measured by the traffic and weather sensors was transferred through a fibre optic communication network to

the processing servers. The processing servers received the time-based (weather, traffic) and event-based data (extended floating car data, xFCD) and calculated a recommended adaptive velocity for each lane and section along the track. The algorithms developed for the calculation of the adaptive velocity were also a subject of the research project.

The aim of the project was to demonstrate traffic safety applications that run within a cooperative vehicle-infrastructure system. A number of test scenarios were run during the project in which incidents on the infrastructure side were simulated by manually changing the traffic or weather values. The data transmission time between the infrastructure and vehicles was under two seconds, which is a satisfying value at a prototype (demonstration) level. The uploading of extended floating car data events was also demonstrated successfully by sending (simulated) in-vehicle events (eg 'accident'). The information

was processed and sent back to the relevant section (and vehicles) within two seconds.

At the beginning of the project, an accident analysis for this specific part of the freeway was performed. This showed that almost 50% of accidents recorded were due to cars straying out of their lanes due to unsuitable velocities, while another 20% were rear-end collisions. Both these types of accident can be avoided by a warning and recommendation system, and a significant reduction in the number of accidents and their severity can be expected in the future.

Link:

<http://www.arsenal.ac.at>

Please contact:

Kashif Din

Österreichisches Forschungs- und
Prüfzentrum Arsenal GmbH, Austria

Tel: +43 50550 6639

E-mail: kashif.din@arsenal.ac.at

Realizing Ambient Assisted Living Spaces with the PERSONA Platform

by Francesco Furfari and Mohammad-Reza Tazari

The PERSONA project is working on a general-purpose technological platform for developing and deploying Ambient Assisted Living (AAL) services. Initial results include middleware that supports self-organization in an ad-hoc manner and a reference architecture for building AAL spaces.

The PERSONA project (Perceptive Spaces prOmoting iNdependent Aging) is an EU-funded research project (FP6) that commenced in 2007. It aims at developing a scalable open standard technological platform for building a broad range of Ambient Assisted Living (AAL) services. AAL is the concept that embraces all those technological challenges involved in addressing the problem of population aging in Europe.

The main technical challenge for PERSONA is the design of a self-organizing middleware infrastructure that guarantees the extensibility of component/device ensembles in an ad-hoc fashion. In order to achieve this goal, the communication patterns of the infrastructure must be based on distributed coordination strategies for service discovery and utilization.

The components of a PERSONA system interact using the PERSONA middleware, which allocates a number of communication buses, each adopting specific and extensible communication strategies. Components linked to the PERSONA middleware may register with some of these communication buses; using P2P connections between the middleware instances, instances of the buses collaborate to enable interoperability among components. Currently four types of bus cover the high-level communication needs in AAL spaces: the input bus, the output bus, the context bus and the services bus.

PERSONA uses a connector-based approach to implement an extensible communication mechanism between distributed instances of the middleware (peers). The current prototype of the PERSONA middleware implemented on

the OSGi platform (see Figure 1) uses connectors based on UPnP, Bluetooth and R-OSGi technologies.

Figure 2 outlines the conceptual architecture of PERSONA with its bus system and shows some of the platform components necessary for constructing AAL services. They are derived from a thorough analysis of different PERSONA scenarios in a set of crucial sectors: social integration, daily activities, safety and mobility (see <http://www.aal-persona.org/scenarios.html>).

As can be seen in Figure 2, the multimodal input processing components that register with the input bus capture user utterances and publish their interpretation of explicit user input to the system. Similarly, output-processing components register with the output bus in order to present system output to the

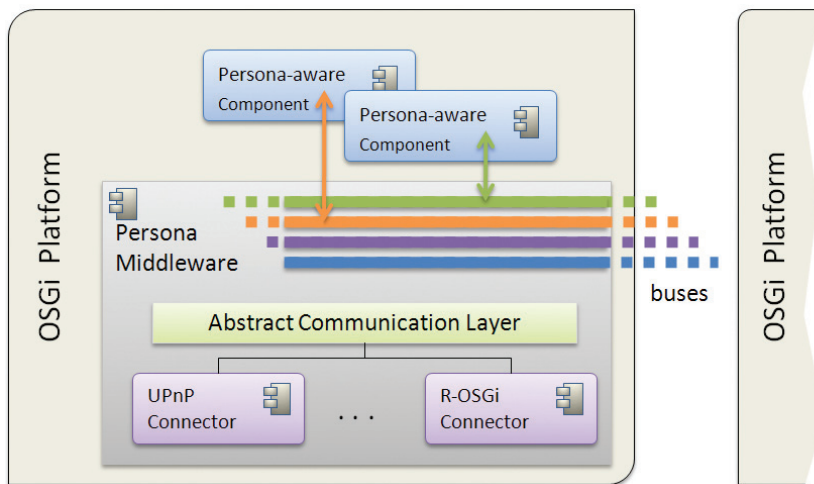


Figure 1: Instances of PERSONA middleware (peers) connected by virtual buses.

user. The context bus is also an event-based bus, to which context sources are attached; the events generated may be reprocessed and transformed into high-level events (situations) by components subscribed to the bus (eg context reasoners). The service bus facilitates access to all the services available in AAL spaces. These services can be atomic or composite; the latter are managed by a Service Orchestrator component. Services registered to the service bus may be requested by other services and by the Dialog Manager in consequence of user input or contextual/situational events according to rules stored in the Knowledge Base of the system.

In order to address issues of interoperability and to add reasoning capabilities, PERSONA has chosen an ontological approach for user and context modelling, defining composite services and adaptation rules, and designing the user interactions with the Dialog Manager. The consortium is currently working on the specification of platform components such as the Dialog Manager, along with the definition of domain ontologies that extend the upper ontologies provided by each of the communication buses. Developments related to the PERSONA platform are planned to be open source.

As is generally known, an ontology provides a formal and explicit model of a specific domain, which can be used to make assertions related to the domain resources and their properties. The information provided by the model and

the known facts based on it form a knowledge base that, in the case of AAL spaces for instance, can serve as extensional information for AAL space configuration.

As shown in Figure 2, PERSONA is made highly configurable by the externalization of metadata describing situations and services as well as the rules associating service invocations with situational events. A problem, however, is that the pre-configured knowledge with which AAL spaces may be equipped is

generally subject to change according to user requirements and/or environment layout. It is therefore important that users be able to directly adapt the knowledge base. However, end users may not be able to cope with the complexity of ontological formalism (eg OWL DL). In the semantic Web and the field of Web page annotation, ongoing research is therefore trying to bridge this gap by providing editors that use controlled natural languages and offer guided user interfaces. The PERSONA platform may adopt a similar approach to annotate AAL spaces formed from its platform components and distributed pluggable context sources and service providers. Combining this approach with other W3C specifications, such as the recent 'Semantic Interpretation for Speech Recognition (SISR)', would provide end users with a scripting facility allowing them to adapt the system according to their specific needs.

Link:
<http://www.aal-persona.org/>

Please contact:
 Francesco Furfari
 ISTI-CNR, Italy
 E-mail: francesco.furfari@isti.cnr.it

Mohammad-Reza Tazari
 Fraunhofer-IGD, Darmstadt, Germany
 E-mail: saied.tazari@igd.fraunhofer.de

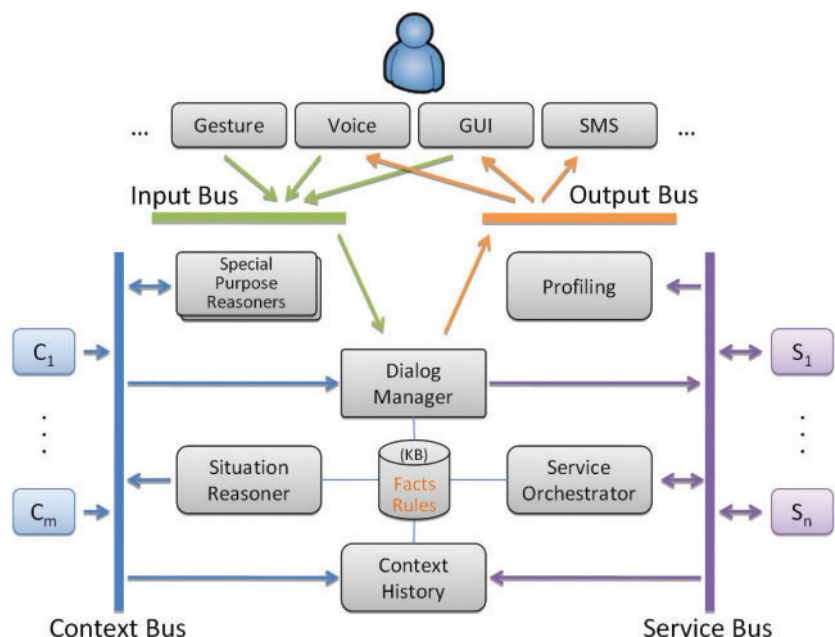


Figure 2: Conceptual architecture of PERSONA.

Medical Record Keeping Made Visually Accessible

by Harry Rudin

A huge medical problem has to do with handling the details of medical treatment – not only keeping the treatment records themselves – but also communicating the treatment to the patient. With a new technique the record keeping is associated with an "avatar", a computer-generated and computer-manipulated 3-dimensional figure representing the patient.

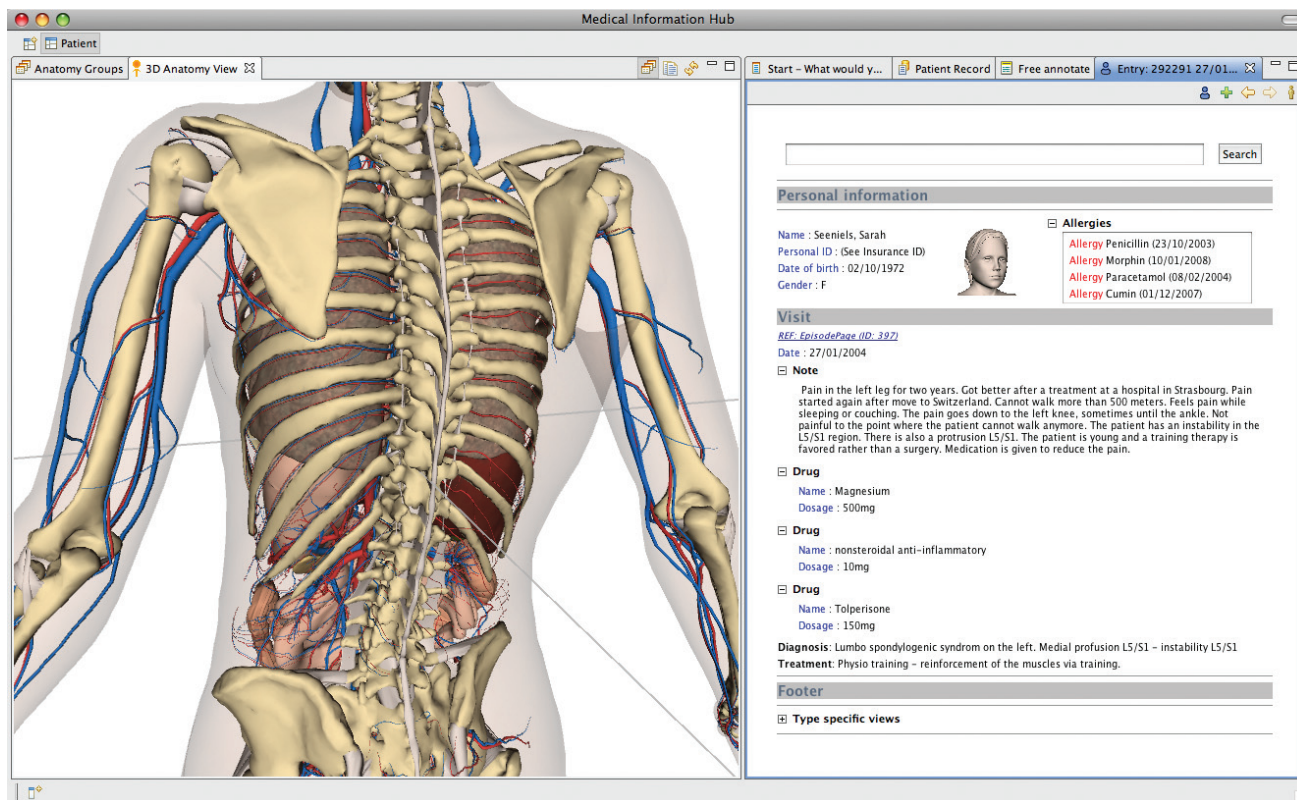


Figure 1: Screenshot of the medical record.

Today, the most comfortable and effective means of handling information is a combination of 3-D graphics, hyperlinks as used in the World Wide Web, and free-text search. Using advanced machine-learning techniques, these approaches are combined in this project initiated at the IBM Zurich Research Laboratory.

The 3-D avatar is a computer-generated figure representing the patient's body. It is easy to display the different parts of the body - the skin, the circulatory system, the skeleton or any of the organs. A mouse click suffices. Medical information about the patient is treated as a personal Web and linked to the avatar via hyperlinks. Textual information about the patient (so-called unstructured data)

can be searched and is automatically linked to the patient's graphical representation through use of SNOMED (Systemized NOMenclature of MEDicine), a collection of some 350,000 standardized medical terms.

The system being developed does have an official name, ASME, which stands for Anatomic and Symbolic Mapper Engine. It was first announced in September of 2007. An informative glimpse of the power of the system can be had through a YouTube video of the system in operation; the URL is given below. Work continues on the development of the system in cooperation with IBM Denmark and a field trial at the Thy-Mors Hospital in northern Denmark.

The purpose of the ASME project is twofold:

- exploring and defining the components of a next-generation browser for electronic medical/health records
- redesigning access to medical information to lower the barrier of IT acceptance in healthcare organizations.

The result is a comfortable, intuitive and inviting interface. One of the lead researchers, Andre Elisseeff, has described the system as "Google Earth for the body."

The general problem of handling medical data is well known and other institutions are also working on solutions. Microsoft, for example, has their

"HealthVault" system, available on the Web in beta test. Google also has a system available in beta test, "Google Health." Both the Microsoft and Google systems are intended for the individual or family user to collect and store health information.

The result of the ASME project is making interaction with medical data as simple and as convenient as possible. If, for example, the doctor were to click on the spine, the system would display

the available medical history, including text entries, lab results and medical images such as radiographs or MRIs. While the system is already extraordinarily attractive, work continues to make it still better. One idea is to include speech technology to ease record keeping.

I look forward to finding such a device in my doctor's office.

Harry Rudin, Swiss Editor, ERCIM News

Link:

Video: <http://www.youtube.com/watch?v=VAwYdmUd59A>

Please contact:

Andre Elisseeff
or Ulf Nielsen

IBM Zurich Research Laboratory,
Switzerland

Tel: +41 44 724 8686

E-mail: ael@zurich.ibm.com

or uln@zurich.ibm.com

Games on Networks

by Jacek Miękiś

Is it always better to be informed? How can the Prisoner's Dilemma be solved? Answers to such questions may be provided by game theory. A team at the Institute of Applied Mathematics and Mechanics at the University of Warsaw explores the unexpected consequences of being in a Nash equilibrium. But is this equilibrium really desirable?

Is it worth trying to keep oneself informed? The answer seems obvious to any rational person, but is it really so? Imagine two companies making the same product. In setting their production levels, they face two conflicting factors: they would like to increase production in order to sell more of the product, but on the other hand, a larger supply means a lower price per unit.

Let us suppose that the more aggressive company announces its production level. We may think that the second company is then in a better position because it is able to adjust its production accordingly. Wrong! Moreover, it is actually better for the second company to avoid observing the production levels of the first, and in addition to make sure that the first one knows this.

How this can be true? Can we justify or prove these statements? To do so, we must translate the problem into mathematics: that's when game theory enters. We have two players, with their strategies being production levels and their payoffs as profits (revenue from product sales less production costs). The question is what they should do. How should they play? The fundamental

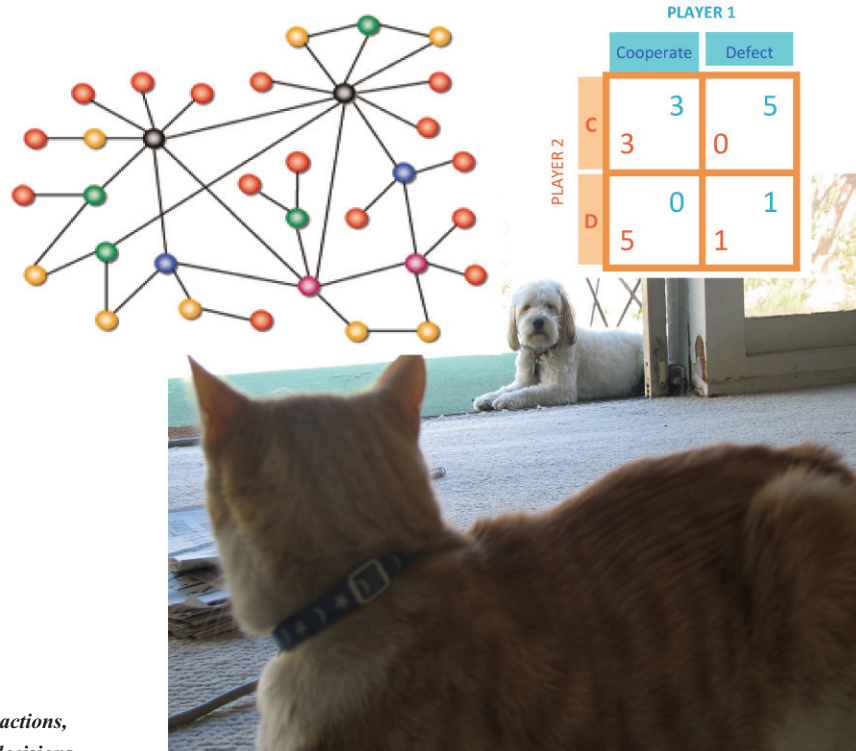


Figure 1: Graph of social interactions, Prisoner's Dilemma, and real decisions.

concept is the Nash equilibrium - the assignment of strategies to players such that no one player can improve his payoff by changing only his own strategy. In the Nash equilibrium of the above game (the classic Stackelberg Duopoly), the profit of the first company is greater than that of the second, but the profit of the second company is increased if it does not know the production level of the first. Thus, it is better to be the first player, but if you are the second one you are better off not being informed.

There are many other socio-economic and biological processes involving interacting agents that can be modelled by game theory. Some of these problems have counterintuitive solutions, making it necessary to solve them mathematically or by numerical simulations rather than attempting to reason them out.

Cooperation between unrelated individuals in animal and human societies is an intriguing issue in biology and social sciences. This problem is described by the Prisoner's Dilemma, in which two players have two options: to cooperate or to defect. The game has a unique Nash equilibrium where both players defect. In fact, defection is the best response to both cooperation and defection of the opponent. However, both

players are better off when they cooperate (see the payoff matrix in Figure 1).

How can this social dilemma be solved? One solution is to interact not with everybody, but only with neighbours in your social network (which might look like the one in Figure 1). In this way many groups of linked cooperators are formed and are difficult to destroy. In spatial games, players are located on vertices of certain graphs and interact only with their neighbours. They may update their strategies using different dynamical rules such as best-reply or imitation. It has recently been observed that real social networks are of certain random types, the most popular being the Barabasi-Albert scale-free graph grown by a preferential attachment rule – the greater the number of people to whom you are connected, the more likely you are to acquire a new colleague.

In a current project, 'Mathematical models of evolutionary game theory and genetic networks', supported by the grant from the Polish Ministry of Science and Higher Education and conducted in the Institute of Applied Mathematics, University of Warsaw, we study the long-run behaviour of spatial games, in particular the Prisoner's Dilemma on different networks. We observed for example that the

density of cooperators undergoes abrupt changes with respect to various parameters in our models.

In many situations, multiple equilibria exist. We address the fundamental problem of equilibrium selection and study equilibrium transitions, using for example the methods and techniques of statistical physics. We explore the similarities and differences between socio-economic systems of interacting players and physical systems of interacting particles.

Networks themselves are formed as a result of the strategic choices of players. Moreover, in games on dynamic networks, agents may simultaneously choose their actions and co-players. We plan to investigate different approaches to the co-evolution of network structures and player strategies.

Link:
<http://www.mimuw.edu.pl/~miekisz>

Please contact:
Jacek Miękiś
Institute of Applied Mathematics
and Mechanics, University of Warsaw,
Poland
Tel: +48 22 5544469
E-mail: miekisz@mimuw.edu.pl

Carrot²: Making Sense of the Haystack

by Stanisław Osiński and Dawid Weiss

Searching and making sense of a large collection of text documents has always been a problem. By using tools that can automatically organize and present documents in a structured form, users can improve their productivity and speed up access to information.

Looking for a piece of information in large collections of documents (such as the Internet) is very much like looking for a needle in a haystack. One can limit the number of matching documents returned from a search engine with the right choice of keywords, but this usually requires some initial knowledge of the context in which the information in question may appear. When this context is unclear, or if this very context is the resource sought, information retrieval becomes a great challenge.

Carrot² is a collection of algorithms and tools designed to help humans explore the thematic context of documents

retrieved from a text collection. A set of documents (for example a list of results retrieved from a search engine) is analysed and dynamically linked together into groups related to a common topic (see Figure 1). Typical examples used to demonstrate this technique show the context of broad and ambiguous queries like 'apache' (helicopter, indian tribe or software organization) or 'salsa' (dance or food). However, the gains from explicit visualization of context in information retrieval go far beyond simple queries. Companies such as Google, Amazon and Vivísimo already employ techniques for context exploration and visualization to

improve their search products. Carrot² offers best-of-breed algorithms and contains demonstration applications for clustering data from multiple sources, including search engines like Google, Yahoo! and MSN, and data repositories like Wikipedia or PubMed.

Carrot² was established in 2001 by Dawid Weiss and Stanisław Osiński, who at the time were students at Poznań University of Technology in Poznań, Poland. From its inception the project was meant to be open and provide value to both the research and commercial communities (BSD licence). The source code was published at SourceForge and

further development took place in the public domain.

From a research point of view, the task of text clustering presents a great challenge, especially in a multilingual context. While a number of document-clustering techniques exist, they all lack the fundamental ability to provide sensible descriptions (labels) of the output docu-

ment groups. This has been the primary focus of the Carrot² project – to extract sensible groups of documents on related topics, but most of all to provide a short, comprehensive description of these clusters.

At the time of writing, the project includes a number of original text-clustering algorithms and auxiliary compo-

nents for text processing. In 2004, Carrot² was awarded a special prize for research tools in the finals of the European Academic Software Award competition. The rough-set-based clustering algorithm included in the project received best paper award at the 2005 IEEE Web Intelligence conference. We are also proud to have a number of deployments worldwide, many references in research literature and a few sibling open-source projects using Carrot² components. Constantly growing commercial interest in text-clustering services and algorithms resulted in the establishment in 2005 of a spin-off company, Carrot Search. The company took over the maintenance and further development of the project.

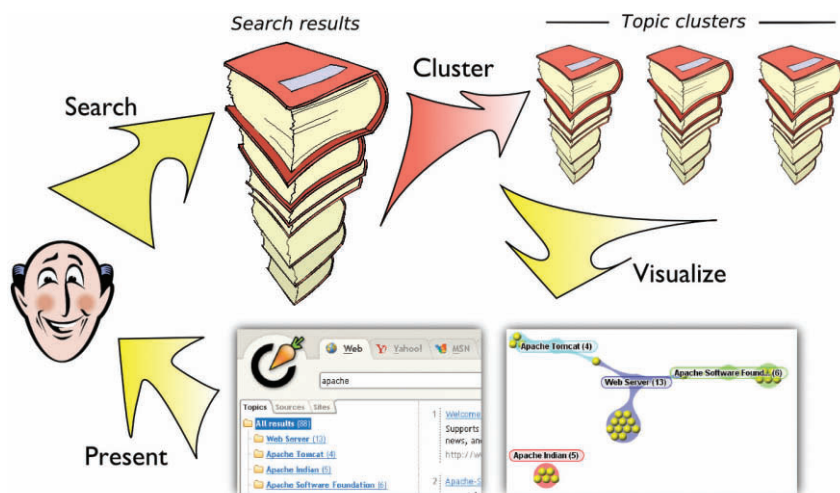


Figure 1: Information flow inside Carrot2 – a set of search results is clustered into topic groups and then shown back to the user in a variety of ways (hierarchy of topics, graph of relationships, etc).

Links:

<http://www.carrot2.org>

<http://www.carrot-search.com>

Please contact:

Dawid Weiss

Institute of Computer Science, Poznan
University of Technology, Poland

E-mail: dawid.weiss@cs.put.poznan.pl

Quantifying WiMAX Performance

by Kostas Pentikousis, Ilkka Harjula, Esa Piri and Jarno Pinola

According to some estimates, by 2010 WiMAX operators will cover areas inhabited by more than 650 million people. Yet current deployment lags similar previous expectations and predictions. Moreover, non-vendor, third-party empirical evaluations of the technology are far from common. We recently concluded our VoIP and IPTV synthetic traffic measurement and analysis studies over fixed WiMAX at the VTT Converging Networks Laboratory in Oulu, Finland, and briefly report our results here. Moreover, we review our simulation-based studies of extending WiMAX in demanding, mountainous terrains, which are not considered to be the main target deployment environments.

WiMAX, often cited as technology that could serve as the substrate for next-generation mobile broadband networks, stands for Worldwide Interoperability for Microwave Access and is based on IEEE 802.16 standards. WiMAX networks can provide point-to-point and point-to-multipoint broadband IP connectivity to both fixed and mobile hosts, with Quality of Service (QoS) guarantees and robust security.

In theory, and according to vendor field trials and demonstrations, WiMAX can

deliver cell bitrates greater than 100 Mb/s, covering large areas (up to 50 km radius from a single base station site using directional antennae), and serving tens of subscribers. These are impressive figures. However, the currently available, commercial off-the-shelf (COTS) equipment delivers significantly less application-layer throughput. At VTT we empirically evaluated fixed WiMAX, aiming to improve our understanding of what is realistically possible using COTS equipment. In particular, we employed the VTT Converging Net-

works Laboratory (CNL) infrastructure, which includes a fixed WiMAX base station (BS) and two subscriber stations (SS), and studied Voice over IP (VoIP) and live IPTV streaming uplink and downlink performance (see Figure 1). Fixed WiMAX was used both as backhaul for voice and data services as well as a last-mile network access technology. CNL is connected via GEANT2 to the Internet.

We employed multiple competing traffic sources over a point-to-multipoint

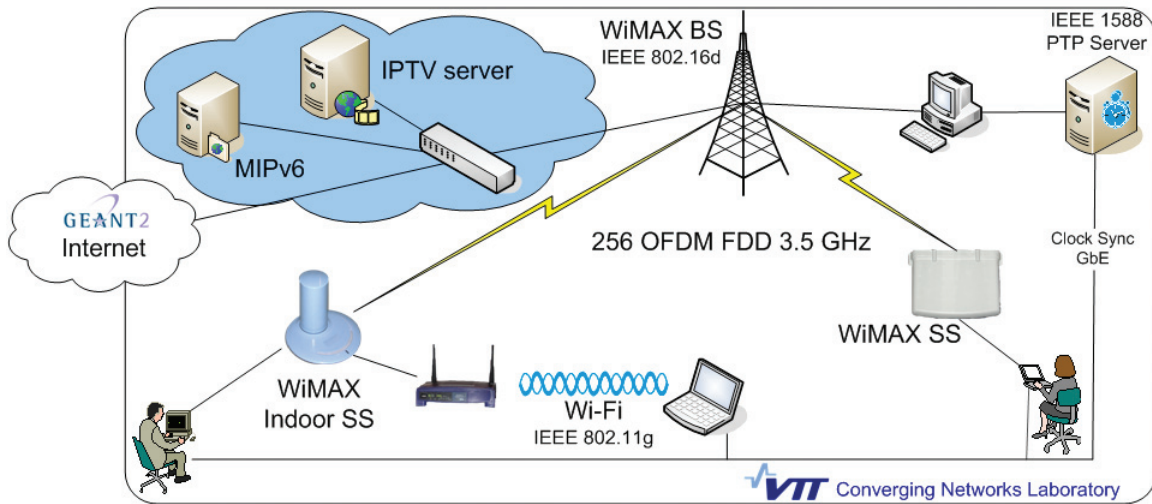


Figure 1: Schematic of the VTT Converging Networks Laboratory fixed WiMAX testbed.

WiMAX topology; its capacity was measured in terms of number of synthetic bidirectional VoIP 'calls' between subscriber stations while concurrently delivering a variable number of video streams with negligible loss. For VoIP, we considered several scenarios using both Speex and G.723.1 codecs. The video stream was captured from a live TV channel transmission and retransmitted in H.264/AVC format.

We measured throughput, packet loss and one-way delay for both line-of-sight (LOS) and non-line-of-sight (NLOS) conditions. We also calculated mean opinion scores (MOS) based on the ITU Telecommunication Standardization

Sector (ITU-T) E-model, for the experiments with G.723.1. We accurately measured one-way delay by employing a software-only implementation of the IEEE 1588 Precision Timing Protocol (PTP). Finally, in order to put our results in perspective, we repeated the measurements after adding IEEE 802.11g access points (APs) in the topology of Figure 1. We found that VoIP flows carrying single sample payloads are clearly underperforming and that application-layer VoIP aggregation can more than triple the number of lossless VoIP flows in the downlink without any network or hardware support. Moreover, our results indicate that a single WiMAX subscriber station can

backhaul VoIP traffic from at least two Wi-Fi APs. Due to space constraints we cannot cover all our results in this article, and refer readers to the peer-reviewed publications available from the links below.

Besides empirically quantifying COTS WiMAX performance, we also used Monte Carlo simulations to study the performance of the WiMAX physical layer in more challenging radio environments. Radio channel properties in isolated mountainous areas differ significantly from those considered in the main target environments of the WiMAX system specifications. We therefore studied promising ways of compensating for performance loss due to environmental factors by steering the direction of the receiver antenna beams. First, the radio channel properties in the mountainous environment around the Vesuvius volcano in southern Italy were analytically derived based on the bistatic radar equation and the geological properties of the area. Careful analysis revealed that the presence of the mountain increases the length of the channel delay spread significantly as well as spreading the received signal in the spatial domain. This model was merged with the WINNER I channel model in Matlab in order to generate a model that could be used in computer simulations.

Consequently we studied the ability of several MIMO algorithms, especially the so-called pre- and post-FFT EVD (Fast Fourier Transform Eigenvalue Decomposition) beamformers, to compensate for performance loss caused by

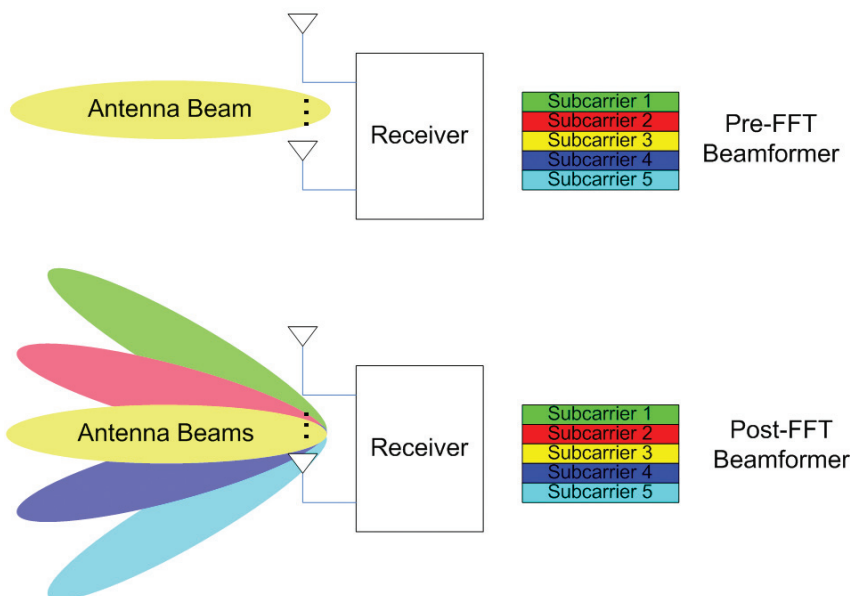


Figure 2: Pre-FFT (top) and post-FFT (bottom) beamformers.

the long channel impulse response. While a pre-FFT EVD beamformer steers the whole receiving antenna beam to one direction, a post-FFT beamformer behaves in a frequency-selective manner and steers the receiver antenna beam for each Orthogonal Frequency Division Multiplexing (OFDM) subcarrier separately (see Figure 2). The simulation results demonstrated that both types of beamformer were able to compensate for the performance loss caused by the mountainous envi-

ronment. However, the post-FFT EVD beamformer demonstrated superior performance at the cost of higher computational complexity and sensitivity towards implementation imperfections, such as inaccurate channel estimates.

The work described herein was mainly undertaken in the framework of WEIRD, a European 'Integrated Project' funded by the 6th Framework IST programme, which was briefly introduced in ERCIM News no. 73.

Links:

VTT Converging Networks Laboratory:
<http://cnl.willab.fi>

IST WEIRD: <http://www.ist-weird.eu>

Related publications:

<http://ipv6.willab.fi/kostas/?weird>

Please contact:

Kostas Pentikousis

VTT Technical Research Centre
of Finland

Tel: +358 40 536 9052

E-mail: kostas.pentikousis@vtt.fi

DEPLOY: Industrial Deployment of Advanced System Engineering Methods for High Productivity and Dependability

by Alexander Romanovsky

The work of the FP7 ICT DEPLOY Integrated Project is driven by the tasks of achieving and evaluating industrial take-up, initially by DEPLOY industrial partners, of DEPLOY methods and tools, together with the necessary further research on methods and tools.

Formal engineering methods enable greater mastery of complexity than do traditional software engineering processes. It is the central role played by mechanically-analysed formal models throughout system development that enables mastery of complexity. As well as leading to big improvements in system dependability, greater mastery of complexity leads to greater productivity by reducing the expensive test-debug-rework cycle and by facilitating increased reuse of software.

The successful three-year FP6 STREP RODIN project on Rigorous Open Development Environment for Complex Systems (2004-2007, <http://rodin.cs.ncl.ac.uk/>) researched and developed industrial strength methods and tools paving the way for the technology to be deployed. In particular, RODIN delivered an extensible open source platform, based on Eclipse, for refinement-based formal methods along with a body of work on formal methods for dependable systems. DEPLOY exploits and builds on these results.

In DEPLOY five leading European companies, representing five major sectors: transportation (Siemens), automotive (Bosch), space (Space Systems), telecommunication (Nokia) and business

information (SAP), will deploy advanced engineering approaches to further strengthen their development processes in order to improve competitiveness

Objectives

The overall aim of DEPLOY is to make major advances in engineering methods for dependable systems through the deployment of formal engineering methods. The work is driven by the tasks of achieving and evaluating the industrial take-up of the DEPLOY methods and tools, initially in the five sectors which are key to European industry and society.

The aim will be achieved with a coherent integration of scientific research, technology development and industrial deployment of the technology. The complementary expertise and technological base of the industrial deployment partners and the technology provider partners will be combined to achieve a set of challenging scientific and technological objectives.

Consortium

DEPLOY offers a balanced interplay between industrial deployment, scientific research and tool development, where companies in five sectors join their forces with eight technology providers to meet the goal.

The industrial sectors, transportation (Siemens), automotive (Bosch), space (Space Systems), telecommunication (Nokia) and business information (SAP), comprise a palette of important European base industries of today. The companies possess different maturity levels when it comes to deploying formal approaches.

The five academic partners are world leaders in formal methods research, that have considerable experience in developing and applying resilience methods as well as a wide range of formal approaches.

The tool vendors, Systerel and ClearSy, have long-standing experience in developing tool support for formal engineering methods. CETIC has considerable experience in industrial quality measurement and will be in charge of the assessment activities.

The project is coordinated by Newcastle University with a dedicated Project Office set at the School of Computing Science. Project partners are Newcastle University (UK), Aabo Akademi University (Finland), Bosch (Germany), CETIC (Belgium), ClearSy (France), ETH Zurich (Switzerland), Heinrich-Heine Universität Düsseldorf (Ger-

many), Nokia Research Center (Finland), SAP AG (Germany), Siemens Transportation Systems (France), Space Systems Finland (Finland), Systerel (France), University of Southampton (UK). Subcontractors are Martyn Thomas (UK) and RATP (France).

Strategies for Deployment

The key to achieving successful and cost-effective deployment of advanced engineering technology will be the construction of strategies for deployment. DEPLOY will develop this concept taking into account the specific characteristics of ways in which rigorous engineer-

- have been field-tested in and adapted for a range of industrial engineering processes
- are accompanied by deployment strategies for a range of industrial sectors
- are based on an open platform (Eclipse) and will themselves be open.

Measurable Outcomes

By the end of DEPLOY each industrial partner will achieve real deployment of formal engineering methods and tools in development of products and become self sufficient in the use of formal engi-

open development platform based on Eclipse that provides powerful modelling and analysis capabilities, is highly usable by practising engineers and is tailored to sector-specific engineering needs. Through the experience and insights gained in the industrial deployments DEPLOY will deliver strategies that enable the integration of formal methods and tools with existing sector-specific development processes.

DEPLOY will put in place an organisation which will be the home of the open platform, set up a body made of industrial users and technology providers whose role

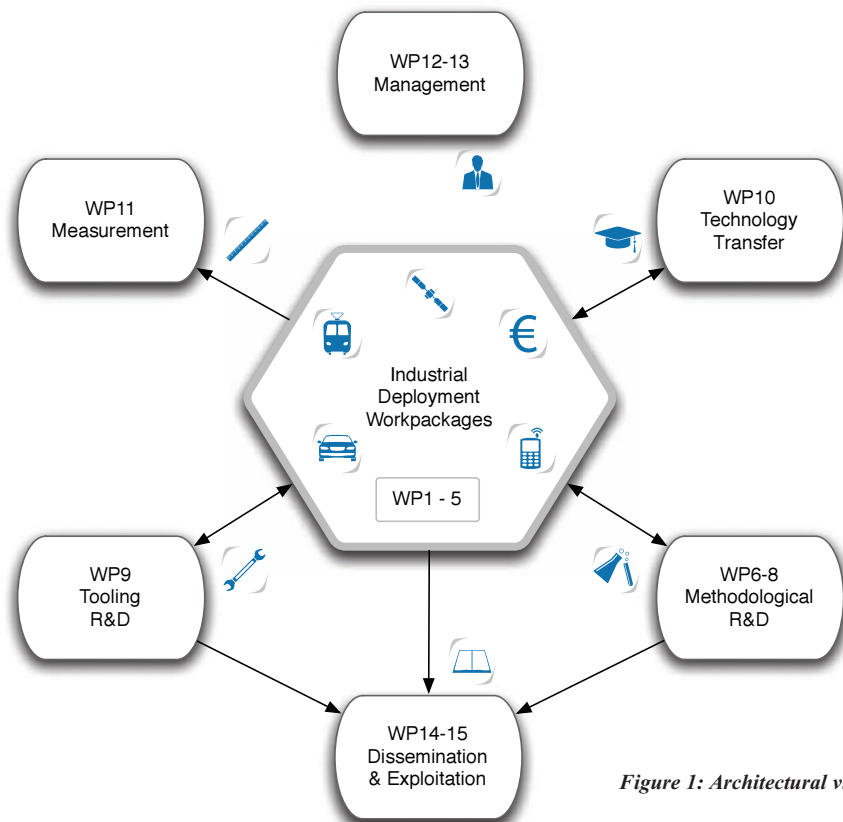


Figure 1: Architectural view of DEPLOY.

ing technologies are applied in each industrial sector.

Results

DEPLOY will deliver methods and tools that:

- support the rigorous engineering of complex resilient systems from high level requirements down to software implementations via specification, architecture and detailed designs
- support the systematic reuse and adaptation of models and software thus addressing industry's requirement for high productivity and requirements evolution

neering methods. The deployments will enable us to provide scientifically valuable artefacts including formally developed dependable systems and results of systems analysis including a rich repository of models, proofs and other analysis results.

By extending the mathematical foundations of formal methods the project will deliver research advances in complex systems engineering methods that enable high degrees of reuse and dependability, and effective systems evolution that maintains dependability. DEPLOY will deliver a professional

will be to coordinate technical decisions on the open platform and deliver training material covering general and sector-specific formal engineering methods.

Link:
<http://deploy-project.eu/>

Please contact:
 Alexander Romanovsky
 Newcastle University, UK
 Tel: +44 191 222 81 35
 E-mail:
alexander.romanovsky@ncl.ac.uk

LREC 2008 - The Language Resources and Evaluation Conference

by Nicoletta Calzolari and Khalid Choukri

LREC, launched in 1998 in Granada as a visionary idea of Antonio Zampolli and other colleagues, celebrated its tenth birthday in the wonderful city of Marrakech. The conference had the honour of receiving the Royal Patronage of His Majesty Mohammed VI, King of Morocco, highlighting the importance assigned to multilingual and multicultural issues.

This year LREC broke many records: we accepted 645 papers out of the more than 900 submissions, and we had over 1,100 participants from all continents. In addition, over 30 satellite events such as tutorials and workshops took place, increasing the chances of discussion and exchange. These figures indicate the great vitality of the broad area of Language Resources (LR) and Evaluation and bear witness to the existence of a strong LR community. The conference was organised by ELRA (the European Language Resources Association) in cooperation with a wide range of international associations and organisations.

Since its creation in February 1995, ELRA (the European Language Resources Association) has focused its activities on a central point of interest: Language Resources and Evaluation. A main motivation for this orientation is to be able to focus on the need for common sharing and utilization of the LRs that are required for research and development activities in the Human Language Technology communities.

LREC has become an important observation post for feeling the pulse of today's initiatives in the field, with the possibility not just of listening to the "best or most innovative" method or technique, but of examining the wide diversity of approaches, the variety of resources for many languages including "minority" ones, new emerging trends, large proj-



LREC Opening Ceremony. From left to right: Bente Maegaard, ELRA President; Taib Debbag, representing M. Ahmed Reda Chami, Moroccan Minister of Trade, Industry and New Technologies; Nicoletta Calzolari, Conference Chair, Khalid Choukri, Chief Executive Officer of ELRA; M. Jochen Richer, representing M. Leonard Orban, the European Commissioner for Multilingualism, Abdelhak Mouradi, co-chair of the local organizing committee.

ects, initiatives and infrastructures. This broadness of themes, topics and perspectives is an essential contribution towards forming a better global vision of our field and thus stimulating new ideas.

We no longer have to defend or promote the 'data-driven' approaches: they are pervasive and have a well-deserved and ample recognition as the necessary infrastructure underlying Language Technology (LT). It is the merit (mostly) of LRs that LT acquires the maturity and attains the robustness needed to be truly usable in applications with a great impact in the society. Nevertheless, LREC does not exclude presentation of other approaches.

The achievement of a worldwide linguistic infrastructure, however, requires the coverage not only of a range of technical aspects, but also – and maybe most critically – of a number of organisational aspects. The consistent growth of the field brings in itself some sort of revolution, and the need for convergence of the many communities now acting separately, such as LR and LT developers, text, speech and multimodal specialists, terminology, semantic web and ontology experts, content providers, linguists and so on. This is one of the challenges for the immediate future, for a usable and useful 'language' scenario in the global and multilingual network. LREC plays an important role in the integration of these various communities.

This does not mean that there are no infrastructural issues still to be discussed and solved. Important examples are the problem of LR maintenance, or the critical issues of interoperability and sharing. The fact that some interesting debates took place during different panels highlights these aspects (eg, panel on data centre missions). The huge growth of the field should be complemented by a reflection on priorities and future strategies. The fact that a number of strategic-infrastructure initiatives are now being launched in all the continents represents an enormous achievement and opportunity for our field. It is a sign that funding agencies recognise the strategic value of our work and the importance of helping coherent growth through a number of coordinated actions. LREC, together with its workshops, is the place where these – and future – initiatives are presented, discussed, and promoted.

All the LREC Proceedings, together with the proceedings of the accompanying workshops, will be made available on the web as a service to the community.

Links:

<http://www.lrec-conf.org/>

<http://www.elra.info>

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

Please contact:

Nicoletta Calzolari

ILC-CNR, Italy

E-mail: nicoletta.calzolari@ilc.cnr.it

Khalid Choukri

ELRA/ELDA

E-mail: choukri@elda.org



Participants of the Delos Summer School 2008.

DELOS Summer School on Preservation

by **Vittore Casarosa**

Another very successful DELOS Summer School on Preservation was held in Tirrenia, Pisa, on June 8-14, 2008. The event was organized by the DELOS Association on Digital Libraries and by the Humanities Advanced Technology and Information Institute (HATII) at the University of Glasgow, with the sponsorship of the UK Digital Curation Centre (DCC), the European project 'Digital Preservation Europe' (DPE), and the Feltrinelli Foundation. This was the fourth edition in this series of summer schools with earlier ones held at Sofia Antipolis (2005), San Miniato (2006), and Pisa (2007).

The school comprised eight three-hour sessions and a full-day workshop on state-of-the-art in preservation. Both the sessions and the workshop were conducted by the distinguished invited lecturers:

- Priscilla Caplan, Florida Center for Library Automation
- Michael Day, UKOLN
- Maria Guercio, University of Urbino & CASPAR
- Hans Hofman, National Archives of the Netherlands
- Perla Innocenti, HATII at the University of Glasgow
- Andreas Rauber, Vienna University of Technology
- Seamus Ross, HATII at the University of Glasgow
- Manfred Thaller, University of Cologne
- Vilma Karvelyte, Faculty of Communication Vilnius University
- Chiara Daniele, FELTRINELLI Foundation
- Beate Sturm, SUB Goettingen
- Sven Vlaeminck, SUB Goettingen
- Helen Tibbo, University of North Carolina at Chapel Hill
- Sarah Jones, HATII and DCC at the University of Glasgow

The main objective of the school was to provide the students with an appreciation of the issues surrounding digital preser-

vation within the context of digital library development and management, and to provide an understanding of the core research in the area of digital curation and preservation. Complete details about the program can be found on the school Web site

The school was attended by 47 students and 14 lecturers, from 20 different countries. For the first time the school was able to support a number of full bursaries, that covered the registration fee as well as the meals and lodging for the duration of the school. The bursaries were granted on a first come first serve basis to students having a reference letter from their professors. That has allowed the participation to the school of a number of very young researchers, who hopefully will contribute in the future to the much needed advancements of the research in the field of preservation.

The interest shown for the school, also on the part of those that in the end were not able to participate, and the enthusiasm of those that actually participated were such that it really seems that the school is satisfying a need in the preservation arena. It is already planned to have a fifth edition of the school in 2009, on June 7-13. Most probably the venue will be the same, as Hotel Continental in Tirrenia, near Pisa, has proven an excellent location, combining the facilities of a modern conference center with the amenities of a beach resort (taken advantage of by the participants during the lunch break).

Link:
Summer School Web site: <http://www.dpc.delos.info/ss08/>

Please contact:
Seamus Ross
HATII at the University of Glasgow
E-mail: S.Ross@hatii.arts.gla.ac.uk

The Fuschl Conversations – A Central Activity of Systems Sciences

by Gerhard Chroust

Once every two years the little village of Fuschl am See, near Salzburg, becomes a centre of the Systems Sciences. During the first full week after Easter in every even year the International Federation for Systems Research (IFSR) invites some 25 to 30 scientists from all over the world to discuss systemic issues and problems in the form of a so-called conversation.


Conversations originated as an alternative to traditional conferences, from the insight that most of the benefit for participants resulted from the discussions and conversations with one another rather than from the formal content or presentations. The initial format for the Fuschl Conversations was built around Social Systems Design, as espoused by Bela H. Banathy. He defined a conversation as:

- a collectively guided disciplined inquiry
- an exploration of issues of social/societal significance
- engaged by scholarly practitioners in self-organized teams
- who selects a theme for their conversation
- which is initiated in the course of a preparation phase that leads to an intensive face-to-face learning phase.

In the Fuschl Conversation by tradition four to five teams with five to seven participants each discuss different, but related topics. The teams share their findings and participants 'visit' other teams.

The 2008 Fuschl Conversation from 29 March to 3 April had 23 participants from eleven countries world-wide. The topics (and thus the teams) were:

- basic concepts of systems sciences



Proceedings of the 2006 Conversations.

Previous proceedings are available at <http://www.ifsr.org/>.

- the trajectory of systems research and practice
- disseminating, accessing and communicating systems knowledge
- quality and excellence in systems research.

Besides the purely scientific endeavour Fuschl also proves to be a deeply human experience freely discussing topics, opinions etc; as one participant said: "When you leave Fuschl, you are a different person."

Initial results of Fuschl will be documented on IFSR's Web site in August 2008. In November 2008 the proceedings will be available in hard-copy form and also on IFSR's Web site. Proceedings of previous Conversations (see Figure 1) can also be found there.

Links:

IFSR: <http://www.ifsr.org>

Fuschl Conversations: <http://ifsr.ocg.at/world/node/45>

Please contact:

Gerhard Chroust

IFSR secretary general and Fuschl 2008 conference chair

J. Kepler University Linz, Austria

E-mail: gc@sea.uni-linz.ac.at

Call for Participation

EACP'08 - International Colloquium "Environnements d'apprentissage pour les communautés de pratique"

Université Aboubekr Belkaid, Tlemcen, Algeria,
18-20 November 2008

The PALETTE project (Pedagogically sustained Adaptive Learning Through the exploitation of Tacit and Explicit knowledge) will organize an international seminar that aims to improve individual and collective learning of communities of practice. Topics include:

- Collaborative Learning items
- methods, tools, evaluation
- creation, usage and sharing of multimedia information
- services, e-services and Web-services, Web 2.0
- development of novel user interfaces
- ontologies.

The seminar language will be French, however, presentations in English are also accepted. The call for papers and posters session will open on 8 September 2008.

The PALETTE project is funded by the European Commission's 6th Framework Programme and managed by ERCIM.

More information:

<http://palette.ercim.org>



www.nis-summer-school.eu

The European Network and Information Security Agency (ENISA) and the Institute of Computer Science of the Foundation for Research and Technology-Hellas (FORTH-ICS) welcome you to the jointly organized Summer School on Network and Information Security. The Summer School's aim is to bring together experts in Information Security, policy members from EU Member States and EU Institutions, decision makers from the industry as well as members of the research and academic community.

committees

STEERING COMMITTEE

- Mr. Andrea Pirotti, Executive Director of ENISA
- Prof. Constantine Stephanidis, Director of FORTH-ICS

TECHNICAL PROGRAMME CHAIRS

- Dr. Sotiris Ioannidis, FORTH-ICS, co-chair
- Dr. Panagiotis Trimintzios, ENISA, co-chair

CONTACT

- Mr. Ioannis Askoxylakis, FORTH-ICS
admin@nis-summer-school.eu



accommodation venue

The Summer School on Network and Information Security (NIS) will take place in Hersonissos, Crete, Greece. The venue of the Summer School is Aldemar Knossos Royal Village (www.aldemarhotels.com). For instructions, please visit the NIS web page (www.nis-summer-school.eu).



ENISA - FORTH Summer School on Network & Information Security



15-19 September 2008, Heraklion, Crete, Greece
Theme: Network Security

keynote speakers

- Lord Toby Harris, House of Lords
- Mr Mikko H. Hypponen, F-Secure

distinguished lecturers

- Dr. Janet Beattie, Glen Abbot Ltd
"Business Continuity Planning"
- Prof. Steven Bellovin, Columbia University
"New Paradigms for Architectural Security"
- Prof. Levente Buttyan, Budapest University of Technology & Economics
"Wireless Networks Security"
- Mr. Ilias Chantzios, SYMANTEC
"Policy Issues related to Network Security"
- Dr. Richard Clayton, Cambridge University
"Economic issues of Network Security"
- Dr. Myriam Dunn, Cavely Center for Security Studies (CSS) ETH Zurich
"Critical (Information) Infrastructure Protection: History, Trends, and Concepts"
- Ms Anne-Marie Eklund-Lowinder, .SE
"DNSSEC: overview, deployment and lessons learned"
- Prof. Mike Fairhurst, University of Kent
"Biometrics & the citizen: broadening the scope of network security"
- Prof. Giusella Finocchiaro, University of Bologna
"Internet Legal Issues"
- Prof. Angelos Keromytis, Columbia University
"Denial of Services Attacks and Resilient Overlay Networks"
- Dr. Achim Klabunde, European Commission DG-INFOS
"Security in the Regulatory Framework for Electronic Communications"
- Prof. Evangelos Markatos, FORTH-ICS
"Emerging Risks for Networks"
- Mr. Jose Marco Muela & Dr. Dimosthenis Ikononou, ENISA
"Resilience of Networks"
- Prof. Bart Preneel, Katholieke Universiteit Leuven
"Network Protocols Security"
- Prof. Vassilis Prevelakis, Technische Universitat Braunschweig
"Lessons Learned from the Vodafone Wiretapping Incident"
- Mr. Marco Thorbruegge, ENISA
"On the Importance of CERT Services"
- SSA West, David A. Jr., FBI Cyber Squad
"Cyber crime in Networked Environments"
- Mr. Ken Van Wyk, Outbreak Security, LLC, Board Member at FIRST and FIRST.org, Inc., and Carnegie Mellon University
"Network Security Incidence Response"
- Prof. Paulo Verissimo, University of Lisbon
"Resilience of Critical Infrastructures"
- Dr. Claire Vishik, Intel
"Security for E-Commerce"



Co-organised by ERCIM

5th Grids@Work Conference

Sophia Antipolis, France, 20-24 October 2008

Further to the success of the IV GRID Plugtests in 2007, ERCIM, INRIA and ETSI organize this year an enhanced event, composed of the 5th GRID Plugtests, during which the "Grids for Finance and Telecommunication Contest" will take place, and a series of conferences and tutorials, including:

- the annual GridCOMP (Effective components for the Grid) conference
- FMCO 2008 conference - Software Technologies Concertation on Formal Methods for Components and Objects"
- European Commission Concertation Meeting: "From Components to Services to Utilities" Units D3 (Software & Service Architectures & Infrastructures) and F3 (eInfrastructures)
- ProActive/GCM Hands-On Tutorial.

The Plugtests event is supported by the EchoGRID project (<http://echogrid.ercim.org>) and will take place on a Grid infrastructure with many different machines and protocols.

The Grid Computing Opportunity

Grid computing offers a model for solving massive computational problems using a large number of computers arranged as clusters, embedded in a distributed telecommunications infrastructure.

Goals of the Event

- bring together Grid researchers, Grid industrials and Grid users
- learn about the future features needed for the Grid infrastructure, through the Grid users experience
- learn how to best program Grid aware applications;
- get useful feedback on the deployment and interoperability of Grid applications on various Grid clusters
- boost the Industry take-up of innovative Grid technologies
- learn more on Grid for industrial applications and interoperability, with a focus on finance and telecommunications.

Who should come to the Event

- researchers and engineers willing to hear about programming techniques for the Grids
- software and hardware vendors, service providers and end-users of various industrial sectors.

Deadline for registration is 22 September 2008

More information:

<http://www.etsi.org/plugtests/GRID2008/GRID.htm>



Call for Participation

2008 CoreGRID Symposium

Las Palmas de Gran Canaria, Spain, 25-26 August 2008

The 2008 CoreGRID Symposium, organized by CoreGRID Network of Excellence, is held in conjunction with the EuroPar 2008 conference, an annual series of international conferences dedicated to the promotion and advancement of all aspects of parallel and distributed computing.

Grids are revolutionising computing as profoundly as e-mail and the Web revolutionised communications and publishing in the last 15 years. Grids are a crucial enabling technology for Scientific and Industrial development in the decade. Grid, Peer-to-Peer and other types of networked distributed computing and storage technologies, as well as emerging web technologies, have provided new paradigms exploiting distributed resources.

The CoreGRID Symposium aims at being the premiere European event on Grid computing for the dissemination of the results from European and member states initiatives as well as other international projects in Grid research and technologies. The CoreGRID Symposium will focus on all aspects of Grid computing including service infrastructures and as such will bring together participants from Research and Industry. The symposium will cover topics such as :

- Applications
- Agent-mediated approaches and peer-to-peer technologies
- Dynamic composition and orchestration of ubiquitous Grid services
- Experimental testbeds
- Grid Portals
- Grid Services
- Grids and Pervasive Computing
- Grid Information Systems/Services
- Industrial and Business Applications of Grid technologies
- Knowledge and data management
- Network-centric Grid operating systems
- Problem solving environments
- Programming models
- Resource brokering, management and scheduling
- Resource Virtualization.

The CoreGRID Network of Excellence (NoE) is funded by the European Commission's 6th Framework Program and managed by ERCIM. It aims at strengthening and advancing scientific and technological excellence in the area of Grid and peer-to-peer technologies. The network is operated as a European Research Laboratory and brings together a critical mass of well-established researchers from 46 institutions.

More information:

<http://www.coregrid.net/mambo/content/view/578/378/>

<http://158.109.65.1/europar2008/coregrid.php>



Call for Papers

HCI International 2009

San Diego, USA, 19-24 July 2009

The 13th International Conference on Human-Computer Interaction is held jointly with

- Symposium on Human Interface (Japan) 2009
- 8th International Conference on Engineering Psychology and Cognitive Ergonomics
- 5th International Conference on Universal Access in Human-Computer Interaction
- 3rd International Conference on Virtual and Mixed Reality
- 3rd International Conference on Internationalization, Design and Global Development
- 3rd International Conference on Online Communities and Social Computing
- 5th International Conference on Augmented Cognition
- 2nd International Conference on Digital Human Modeling
- 1st International Conference on Human Centered Design

HCI International 2009, jointly with the affiliated conferences, which are held under one management and one registration, invite you to San Diego, California, USA, to participate and contribute to the international forum for the dissemination and exchange of up-to-date scientific information on theoretical, generic and applied areas of HCI through the following modes of communication: plenary / keynote presentation(s), parallel sessions, poster sessions, tutorials and exhibition. The conference will start with three days of tutorials. Parallel sessions, poster sessions and the exhibition will be held during the last three days of the conference. The conference focuses on the following major thematic areas:

- Ergonomics and Health Aspects of Work with Computers
- Human Interface and the Management of Information
- Human-Computer Interaction
- Engineering Psychology and Cognitive Ergonomics
- Universal Access in Human-Computer Interaction
- Virtual and Mixed Reality
- Internationalization, Design and Global Development
- Online Communities and Social Computing
- Augmented Cognition
- Digital Human Modeling
- Human Centered Design

Proposals for Participation

All submitted abstracts will be peer-reviewed by three independent referees from the International Program Boards.

Deadline for Abstract Receipt

Papers: 20 October 2008

Posters: 23 February 2009

Tutorials: 20 October 2008

The HCI International 2009 Conference Proceedings, comprising the papers to be presented at the conference, will be published by Springer in a multi-volume set in the Lecture Notes in Computer Science (LNCS) and Lecture Notes in Artificial Intelligence (LNAI) series.

More information:

<http://www.hcii2009.org>



2nd Edition of the ERCIM-ETSI Infinity Initiative

Bio ICT - The Heart in the Computer

Sophia Antipolis, France 2-3 October 2008

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

Thanks to the upcoming "Virtual Physiological Human", new environments for predictive, patient specific, more effective and safer healthcare are being developed. Can you imagine, for instance, the benefits of having an 'e-drug' tested in an 'e-organ' of an 'e-patient'? To explore this exciting new field supported by advanced Information and Communications Technologies such as Grid Computing, ERCIM and ETSI (European Telecommunications Standardisation Institute) are pleased to invite you to the 2008 edition of their joint 'Infinity Initiative', which will take place in Sophia Antipolis, France, on 2-3 October 2008.

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

ERCIM and ETSI believe it essential that standardisation be considered early in the R&D cycle in order to facilitate maximum creation and development of home and international markets for European industry..

More information:

http://www.etsi.org/website/newsandevents/2008_infinityinitiativeteaser.aspx

EchoGRID at the Open Grid Forum

Over 500 members of the global grid community gathered at the 23rd Open Grid Forum (OGF23) in Barcelona to network, showcase achievements, map out requirements for distributed computing and push forward standardization. The event helped connect the people of OGF and newcomers from across the globe. OGF and the recently EU-funded OGF-EUROPE project strongly support engagement on a truly global level to ensure that people work together to drive openness and interoperability that is of value to business and science. In this respect, OGF welcomed the participation of three EU projects that are highly active in the development and deployment of distributed computing between Europe, China and India: Bridge, EchoGRID and EU-IndiaGrid. The three projects are a good example of global partnerships bringing not just technical value-add but also playing a central role in strengthening relations between Europe and other parts of the world.

EchoGRID (European and Chinese Co-operation on Grid) is an ERCIM-managed road-mapping project that has established synergies with similar initiatives and projects to map out future research priorities in areas like new programming paradigms, grid architectures, service-oriented architectures, Grid management, virtual organisation, component models, and workflows and business processes. EchoGRID, which is actively engaged in Grid open standards, supported an OGF23 session on roadmaps and strategic research agendas featuring NESSI-Grid, Challengers, CoreGRID, 3S, and NEXOF-RA to define common strategies and future activities within OGF.

An extended report on OGF32 by Stephanie Parker has been published in Grid Today.

<http://www.gridtoday.com/grid/2385903.html>

<http://echogrid.ercim.org/>



Workshop Report: Role of Mobile Technologies in Fostering Social Development

W3C published a report on the June 2008 Workshop on the Role of Mobile Technologies in Fostering Social Development. Participants discussed how numerous available services on mobile phones could help people in underserved regions. Discussion underlined the need for a concerted effort among all the stakeholders (including practitioners, academics, regulators, and mobile industry) to build a shared view of the future of the mobile platform as a tool to bridge the digital divide. The Workshop was jointly organized by W3C and NIC.br, as part of the Digital World Forum (DWF) project led by W3C and managed by ERCIM. DWF is supported by the European Union's 7th Research Framework Programme (FP7).

http://www.w3.org/2008/02/MS4D_WS/exec_summary.html
<http://www.digitalworldforum.eu/>

Roberto Scopigno receives Eurographics "Outstanding Technical Contributions" Award 2008

Roberto Scopigno is the recipient of the Eurographics "Outstanding Technical Contributions" Award 2008 in recognition of his outstanding contributions to multiresolution modeling and rendering, 3D scanning, surface reconstruction, scientific visualization, volume rendering and applications to cultural heritage.



Roberto Scopigno.

Roberto Scopigno's work has had a profound impact on the field of visual computing, and several contributions truly stand out. Of particular importance has been his work on surface simplification, LOD and multiresolution representations for surfaces and volumes. Other significant contributions are his work on speeding up isosurface extraction using interval trees, his algorithm for implicit disambiguation of Marching Cubes, his data structures and algorithms for efficient out-of-core memory management and processing of huge meshes, and his stimulating work on 3D scanning with applications to cultural heritage.

Roberto Scopigno leads the Visual Computing Lab of ISTI-CNR and has created a very successful research group, several of his former students have themselves become productive researchers.

Complete award citation:

http://www.eg.org/about/awards/winners/winner_scopigno

Steven J. Murdoch wins the ERCIM Security and Trust Management Working Group Award

Steven J. Murdoch from University of Cambridge received the best thesis award of the ERCIM Working Group on Security and Trust Management for his excellent Ph.D. thesis "Covert channel vulnerabilities in anonymity systems". Steven's thesis represents an outstanding contribution on the study of privacy and anonymity aspects in real systems and meets all the desired requirements. The evaluation committee particularly favoured its exceptional practical impact.

The ERCIM Security and Trust Management (STM) Working Group has recently established an award for the best Ph.D. thesis in the area of STM in order to increase the visibility of the young researchers in the ERCIM and the wider European scientific community. Applications for the 2008

award were open to all Ph.D. holders who defended their thesis during 2007 in a European university: Sixteen applications were submitted and evaluated by a committee of European experts. The theses were judged on the basis of scientific quality, originality, clarity of presentation, and the potential impact of the results.



Steven J. Murdoch.

Steven's thesis:

<http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-706.html>

Security and Trust Management Working Group:

<http://www.iit.cnr.it/STM-WG/>

INRIA continues to host the ERCIM Office

ERCIM's Board of Directors unanimously accepted INRIA's offer to host the ERCIM Office and reappointed Jérôme Chailloux as Manager for the next three years from June 2008 to May 2011. INRIA has been hosting the ERCIM Office since the beginning, providing the ERCIM Manager and supporting the costs for running the office. Michel Cosnard, CEO and Chairman of INRIA, stated: "As a co-founder of ERCIM nineteen years ago, INRIA is proud that ERCIM is now a key actor of the European research area in Information and Communication Science and Technologies."



Jérôme Chailloux, ERCIM Manager.

Jérôme Chailloux was appointed as the ERCIM Manager in June 2005 and as such he is also the Site Manager of the W3C Europe office hosted by ERCIM. Jérôme started his career at INRIA in 1980 where he became a Research Director before co-founding the company ILOG (the second spin-off of INRIA), a world leader in software components for optimisation, decision aid and visualisation. He then took the position of CIO of the genomics company GENSET, where he led one of Europe's largest teams of bioinformaticians. Until 2000, he was a member of the Coordination Committee for Information and Communication Science and Technologies of the French National Ministry of Education, Research and Technology.

Two ERC Starting Grants for INRIA Researchers

Anne-Marie Kermarrec and Bruno Lévy from INRIA received each the "Starting Independent Research Grant" of the European Research Council (ERC). Anne-Marie Kermarrec, head of the ASAP (As Scalable As Possible - Foundation of Large Scale Dynamic Distributed Systems) research team at INRIA/IRISA Rennes, received the grant for her project GOSSPLE. The expected outcome of GOSSPLE are theoretical advances in large-scale distributed computing, practical advances in large-scale distributed systems, an innovative software to create dynamic autonomous networked systems and the ground for new kind of distributed collaborative applications. GOSSPLE has the potential of radically changing the navigation on the Internet by placing users affinities and preferences at the heart of the search process. Complementing traditional search engines, GOSSPLE will turn search requests into live data to seek the information where it ultimately is: at the user.



Anne-Marie Kermarrec and Bruno Lévy.

Bruno Lévy, head of the ALICE (Geometry and Light) project-team at INRIA Nancy received the grant for his project GOODSHAPE (Numerical Geometric Abstraction: from bits to equations), that involves several fundamental aspects of 3D modelling and computer graphics. GOODSHAPE is taking a new approach to the classic, essential problem of sampling, or the digital representation of objects in a computer. This approach will probably result in a new algorithm that is much more efficient in terms of algorithmic complexity and can be applied to a larger class of problems. Possible applications are for example inverse engineering and oil exploration.

Both grants are in the order of one million euro over five years. Anne-Marie Kermarrec and Bruno Lévy have been selected with 297 other researchers, among 10,000 applications in Europe, all sciences included. ERC Starting Grants aim to support up-and-coming research leaders who are about to establish or consolidate a proper research team and to start conducting independent research in Europe. The scheme targets promising researchers who have the proven potential of becoming independent research leaders. It will support the creation of excellent new research teams and will strengthen others that have been recently created.

<http://www.irisa.fr/activites/new/08/AMK>

<http://inria.loria.fr/european-research-council>

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



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



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



Irish Universities Association
c/o School of Computing, Dublin City University
Glasnevin, Dublin 9, Ireland
<http://ercim.computing.dcu.ie/>



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



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



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



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



Centrum voor Wiskunde en Informatica
Kruislaan 413, NL-1098 SJ Amsterdam,
The Netherlands
<http://www.cwi.nl/>



Science & Technology
Facilities Council

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



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



Spanish Research Consortium for Informatics
and Mathematics c/o Esperanza Marcos, Rey Juan Carlos University,
C/ Tulipan s/n, 28933-Móstoles, Madrid, Spain,
<http://www.sparcim.org/>



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



Swedish Institute of Computer Science
Box 1263,
SE-164 29 Kista, Sweden
<http://www.sics.se/>



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

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



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



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



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



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

Fraunhofer Gesellschaft



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



Technical Research Centre of Finland
PO Box 1000
FIN-02044 VTT, Finland
<http://www.vtt.fi/>

Order Form

If you wish to subscribe to ERCIM News
free of charge

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

Send, fax or email this form to:

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

I wish to subscribe to the

printed edition

online edition (email required)

Name:

Organisation/Company:

Address:

Postal Code:

City:

Country:

E-mail:

Data from this form will be held on a computer database.

By giving your email address, you allow ERCIM to send you email