Number 81, April 2010

ERCIM NEWS

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

Special theme:

Computational Science/Scientific Computing

Simulation & Modelling

for Research and Industry

Also in this issue:

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of 9,000 copies. The printed version of ERCIM News has a production cost of €8 per copy. Subscription is currently available free of charge.

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

Keynote

Scientific computing is today becoming the 'third pillar' of scientific inquiry, alongside the more traditional theory and experimentation pillars. At the same time, tough policy choices in the grand challenge areas of our time like energy, the environment, climate and health, will need to be increasingly grounded in Science in order to be effective and justifiable. The availability of high capability systems and simulation software therefore becomes a condition for both competitive Science and sound policy-making. The ongoing evolution towards an integrated European HPC (high performance computing) and computational science infrastructure from scattered national investments now offers Europe the possibility to take a leading position in scientific computing.

HPC systems exhibiting a very high level of parallelism and attaining performances of one or more peta-flops[1] are already being deployed. This is an engineering achievement in terms of balancing performance increases with acceptable power consumption. On the other hand, efficiently exploiting such systems with well-tuned application codes and tools is still a challenge for computational science. In addition, steeply increasing costs for HPC systems and specialised hardware are expected to result in an increasing use of commodity components and shift added value from hardware to software.

In fact, the cost of acquiring and operating a national leadership-class HPC facility, with the periodic investments in hardware, energy and operations, is becoming too high for most European countries to justify bearing it alone. The cost of tuning codes to exploit the stated increases in computing performance is also not negligible. In this context, a consortium of national supercomputing centres decided to form a partnership to create a European HPC infrastructure service composed of leadership-class HPC systems and centres of excellence in computational sciences with the mandate to procure and operate the systems as well as to support users in addressing Grand Challenge problems. This partnership, code-named PRACE[2], is soon to create an organisation that will provide HPC services with funding from several European countries and from the European Union.

Over the previous years, EU support to HPC took the form of: (i) consolidating a distributed infrastructure of national

Kostas Glinos,
European Commission, DG Information Society and Media
Head of GEANT and e- Infrastructure Unit

HPC facilities[3] and (ii) co-ordinating the access of the scientific and engineering community to the computing slots available via a common resource pool[4]. This 'grid' of HPC systems soon became a success with the European computational science community; the services and competence of the infrastructure provider filled a gap, serving many European researchers without national HPC facilities of sufficient performance or without the best match of system architecture to their codes. A policy group was thus formed that called on the European Union to support the creation of a new research infrastructure focussed on the provision of leadership-class systems at the same level as the US and Japan and access to the facility based on scientific merit[5]. Thanks to the work of this policy group, the European HPC service was included in the 2006 ESFRI Roadmap for Research Infrastructures[6] which, in turn, led to the European Union strongly contributing to these developments via the Capacities programme[7].

This process opened a path for progressive convergence of national and EU policies on the strategy for investments in HPC infrastructure. The aim is that the European HPC infrastructure becomes an integrated distributed facility that enables access to computing resources as required by scientific or engineering applications, supporting innovation and scientific breakthroughs and the attractiveness of Europe for research.

In parallel to these activities, the e-infrastructures part of the EU's framework programme has provided support to other ICT-based infrastructure such as the high-speed network GÉANT that interconnects supercomputing centres and the distributed computing grid of EGEE (and soon EGI). Recent additional investments in scientific software, regional HPC facilities and data e-infrastructures complement and enhance the European HPC infrastructure support.

For the vision of EU leadership in HPC to become reality, a number of policy, research and technical challenges still need to be addressed. First, the combined purchasing power of European countries for leadership-class systems needs to leverage a competitive supply of technologies and systems in Europe. Second, the peta-flops performance has been attained but corresponding programming tools for development and maintenance of scientific computing applications are still in the making and no real infrastructure exists for making scientific software available as a service. Third, HPC technology needs to continuously re-invent itself: the next computing performance barrier, the exa-flops[8], will require a new computing paradigm to emerge in order to efficiently attain this level of performance.

Addressing these challenges effectively depends on national and European policy choices, the existence of an efficient coordination framework for all stakeholders and, of course, financial support from national budgets and from the 7th Framework Programme – and from the 8th Framework from 2013 onwards.

Kostas Glinos

---

1 *A peta-flops is the equivalent to ten to the power of fifteen FLoating point Operations Per Second. The peta-flops barrier is the most recently surpassed performance barrier in HPC systems.*

2 *See also more details on PRACE (Partnership for Advanced Computing in Europe) at http://www.prace-project.eu*

3 *See also for more details the project DEISA (Distributed European Infrastructure for Supercomputing Applications) at http://www.deisa.eu*

4 *The DEISA Extreme Computing Initiative (DECI in short) calls for proposals (see also http://www.deisa.eu/science/deci)*

5 *See also the policy work of the European HPC task force (http://www.hpcineuropetaskforce.eu) and the roadmap of new European research infrastructures produced by the ESFRI in 2006 (http://cordis.europa.eu/esfri/roadmap.htm).*

6 *Available from the site: http://cordis.europa.eu/esfri/incubator.htm*

7 *The Capacities Programme aims to enhance research and innovation capacities throughout Europe and ensure their optimal use. It is one of the specific programs of the Seventh Framework Programme (FP7) for research and development (see also the site: http://cordis.europa.eu/fp7/faq_en.html)*

8 *A exa-flops is the equivalent to ten to the power of eighteen FLoating point Operations Per Second.*

Disclaimer: The views expressed are those of the author and do not necessarily represent the official view of the European Commission on the subject.

R&D AND TECHNOLOGY TRANSFER

EVENTS

BOOKS

IN BRIEF

Tribute to Georges Nissen

INRIA and ERCIM have lost one of its pioneers: Georges Nissen, an iconic figure at INRIA who spent his entire career working for the Institute, passed away on Tuesday 9 March 2010.

Georges Nissen joined INRIA in 1970 as a researcher on Alain Bensoussan's "Systems Theory" project. In 1972, he became Corporate Affairs Manager at LABORIA. He was later appointed to INRIA's management board as Jacques-Louis Lions' deputy. Then, in 1986, he took over as INRIA's International and Industrial Relations Director, playing an essential role in collaborations with the United States and Russia.

ERCIM owes a lot to Georges' personal commitment. He was largely involved in the creation of ERCIM and served in the ERCIM Executive Committee from 1989 - 2001.

Georges Nissen will forever be remembered at INRIA and at ERCIM as one of their foremost contributors and one of their best ambassadors.

"On a working level he was very deeply involved in helping to get ERCIM up and running, says Eckart Bierdümpel, the current chair of the Executive Committee. "Always being an active participant in ERCIM discussions he became one of the most respected and longest-serving members of the Executive Committee. As a member right from the beginning in 1989 he was INRIA's voice and face in the Executive Committee for twelve consecutive years. His passion about ERCIM can nicely be demonstrated with the anecdote that he even created a new salutation and once in a while signed his emails with "ERCIMly, Georges". Even after retirement he involved himself in ERCIM activities whenever he was asked to: for example he participated as a contemporary witness of the beginnings of ERCIM at the 20th anniversary of ERCIM, which took place just last year in Paris.

Georges, the Executive Committee will miss you!"

Alain Bensoussan, former President of INRIA and a founder of ERCM adds: "The importance of the role played by Georges Nissen in the birth of ERCIM can be traced back to the contacts between INRIA and its partners GMD in Germany and CWI in the Netherlands. I can remember how much Georges was appreciated by GMD and CWI. His involvement has certainly been instrumental in the shaping of the ERCIM concept. He naturally became a pillar of the EEIG, and has always remained a big supporter.

Thanks Georges, you deserve our gratitude and what you did for ERCIM will not be forgotten."

ERCIM Working Group Models and Logics for Quantitative Analysis Annual Meeting

The second annual meeting of the MLQA Working Group will take place on July 9 July 2010 as part of FLoC 2010 in Edinburgh (Scotland). The theme of the meeting is: "Static Analysis versus Model Checking: similarities, differences, synergies"

The aim is to create a vibrant event consisting mainly of a number of invited talks that will cover some of the historical developments, survey the links established, establish state-of-the-art, identify the problems still worth pursuing and give a perspective on the implications and (novel) applications that can be foreseen.

The following list of presentations have been confirmed (but titles are tentative):

- Bernhard Steffen (on performing static analysis by doing model checking)
- Flemming Nielson (on performing model checking by doing static analysis)
- Marta Kwiatkowska (on quantitative model checking)
- Joost-Pieter Katoen (on quantitative model checking)
- Marsha Chechik / Arie Gurfinkel (on the use of abstraction in model checking)
- Orna Grumberg (on the use of abstraction in model checking)
- David Monniaux (on operations research in static analysis and model checking)
- Michael Huth (on modelling and analysing control of access and risk)

On top of which we will have:

- MLQA Business Meeting (open to all attendees) about the future of MLQA

It will also be possible for attendees to present their own work, either in a poster session or by short 1-slide presentations; the format has not yet been decided.

Link:
For the latest news please take a look at
http://wiki.ercim.eu/wg/MLQA/index.php/July_2010:_MLQA_meeting_at_FLoC_2010%2C_Edinburgh

Please contact:
Flemming Nielson
MLQA Working Group chair
DTU Informatics/DANAIM, Denmark
E-mail: mlqa@imm.dtu.dk)

Euro-India SPIRIT Project Kick-off

by Florence Pesce

Euro-India SPIRIT (Euro-India Synchronisation of Policy Initiatives & Research and Innovation Trajectory) is a Europe-driven initiative with strong Indian partners supporting the objectives of the call "ICT-2009.9.1: International co-operation support to Information Society policy dialogues and strengthening of international cooperation" in specifically contributing to the definition and development of a European and Indian policy dialogue for sustainable ICT policy and research priorities, thereby aiding in the build-up of the European Research Area. The project was lauched at a kick-off meeting in Pisa, 26-27 January 2010.

The goal of "Euro-India SPIRIT" which has a duration from January 2010 - December 2011, aims to engage the EU and Indian ICT stakeholders at a level where policy formulation pertaining to research can be aligned and supported to identify the priorities of key research stakeholders and constituencies. The outcome is to formulate a mutually-beneficial research and innovation agenda that can be taken up through specific bilateral initiatives.

Euro-India SPIRIT builds on the foundations and achievements of successful projects – BASIC - EuroIndia2004 IS Forum, MONSOON – Euro-India ICT Cooperation Initiative and EUROINDIA which ended in December 2009. In more detail Euro-India SPIRIT seeks to:

- analyse the research dimension of Indian ICT policy, encompassing ICT technologies and applications and user industry and domain implications
- organise and probe stakeholders to identify medium to long term vision and perspectives and reveal complementarities with EU priorities, notably the i2010 strategy and objectives and updates
- update, assist and enrich the policy dialogue process and meetings between the EU and India with periodic reports, surveys and position papers encompassing findings from policy research and stakeholder views dealing with common R&D perspectives, priorities, opportunities and challenges
- organise synchronized stakeholder events and domain experts deliberations to enhance the profile and communicate the utility of the policy dialogue process
- recommend future co-operation initiatives and identify matching counterpart funding to leverage Indian research capability and capacity and to engage in projects in common priority areas
- develop synergies with activities underway in India covering Capacities, People and INCO-NET dimensions.

The project's kick-off meeting took place at the Trust-IT Services offices in Pisa, Italy. Consortium members have been joined by the project officer Mr Alvis Ancans from the European Commision, who has been able to participate fully in the proceedings.

Euro-India SPIRIT will be supported by experts' working groups and an online community in order to identify and validate long term common research perspectives aligned with the Indian ICT research policy and EU priorities. Three working groups addressing technological and policy issues involved in the development of sustainable joint EU-Indian ICT research priorities are being created. The working groups will leverage on experience and results from the EuroIndia project and principal ICT topics from FP7 such as focusing on:

- ICT Addressing Societal challenges, (eGovernment, govt & public services, eHealth, etc.
- AudioVisual, Media & Internet
- Emerging Technologies and Infrastructures (Geant, eInfrastructures etc..).

ICT Addressing Societal challenges, (eGovernment, govt & public services, eHealth, etc.)	*AudioVisual, Media & Internet*	*Emerging Technologies and Infrastructures (Geant, eInfrastructures etc..)*
ICT FP7 Challenge 5 - ICT tools for sustainable and personalised healthcare	ICT FP7 Challenge 1 - Pervasive and trustworthy network and services infrastructure	ICT FP7 Challenge 1 - Pervasive and trustworthy network and services infrastructure
ICT FP7 Challenge 6 - ICT for mobility, environmental sustainability and energy efficiency	ICT FP7 Challenge 4 - Digital libraries and content	eInfrastructures through the Research Infrastructures Work Programme of the FP7 Capacities programme
ICT FP7 Challenge 7 - ICT for independent living, inclusion and participatory governance	International Co-operation	Future & Emerging Technologies

Table 1: The three expert working groups of the EU-India SPIRIT projetct.

Each working group will be composed of six members, three European and three Indian members, ensuring a balance of thematic expertise as well as organisational background (industry, public admin, academia) and co-ordinated by the Euro-India SPIRIT partners.

The Euro-India SPIRIT project is coordinated by ERCIM, with the partners Infra Technologies Sarl, France; Trust-IT Services Ltd. ,United Kingdom; The Federation of Indian Chambers of Commerce & Industry, India;Centre for Science, Development and Media Studies, India; and the Administrative Staff College of India, India.

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

Please contact:
Florence Pesce, ERCIM Office
Tel:+33 4 92 38 50 73
E-mail: florence.pesce@ercim.eu

Call for an ERCIM Working Group on Agent-Based Technology

There are many applications of agent-based technology in the mathematical, physical and social scientific communities. Multi-agent systems are very complex applications to implement and optimise for large populations of agents. The optimisation of agent interactions and the management of initial data and outputs for large scale simulations are but two areas that need research and development efforts. Although there are many conferences where work on agent-based system is presented we would like to find out whether there is a sufficient number of interested groups within the ERCIM community to merit setting up an ERCIM Working Group on Agent-Based Technologies.

There are clearly many areas within agent-based methods that might benefit from such activity: serial and parallel implementation of agent-based applications, verification and validation of agent-based models, standards for the exchange of agent-based models to name a few.

If you have an interest in proposing and participating in an ERCIM Working Group on Agent-Based Technology please contact Prof Chris Greenough of the Software Engineering Group at the STFC's Rutherford Appleton Laboratory with your contact details and areas of interest.

Please contact:
Chris Greenough
STFC Rutherford Appleton Laboratory
Tel: +44 1235 45307
E-mail: christopher.greeenough@stfc.ac.uk

ERCIM
European Research Consortium for Informatics and Mathematics

"Alain Bensoussan"
Fellowship Programme

Application deadline:
Twice per year
30 April and 30 September

Photo by courtesy of INRIA.

Who can apply?
The Fellowships are available 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 ERCIM Fellows:
- 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 cross-fertilisation and cooperation, through the Fellowships, 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.

How to apply?
Applications have to be submitted online. The application form will be available one month prior to the application deadline at http://www.ercim.org/activity/fellows/

Which topics/disciplines?
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 Methods, Grids, Security and Trust Management, and many other areas in which ERCIM institutes are active.

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

New EC-funded Projects granted to ERCIM/W3C

Multilingual Web

Multilingual Web is a thematic network composed of 20 partners, including the ERCIM members CNR, UPM (SpaRCIM) and University of Limerick (IUA). By raising the visibility of standards and best practices for the multilingual Web, the network will set a foundation for improved support on the Web for languages of the European Union and its trade partners. These improve the efficiency of processes for creating and localizing content, both by machine translation and more traditional methods, and improving support for multilingual content and data on the Web. This W3C/ERCIM managed project will officially start in April 2010.

Open Media Web

Web technology is currently undergoing a sea-change which makes it a strong contender for becoming an open, royalty-free standards-based platform for networked media that levels the playing field and enables new market participants, including European players. The Open Media Web project aims to build on Europe's strength in multimedia technology and content to enable European research and industry to strengthen its position in Web technology. The project is led by the W3C and ERCIM. The project started in January 2010.
Link: http://openmediaweb.eu/

Future Research Challenges in Software Evolution and Maintenance – Report from EC Expert Meeting

by Tom Mens

In Fall 2009, on behalf of the EC DG Information Society, Arian Zwegers convened a group of experts on the theme of Software Evolution and Maintenance, in order to discuss about the medium and long term research challenges (timeframe of 2015 and beyond). To the extent possible, the EC will take into account these challenges while preparing their next framework workprogramme.

The expert group was moderated by Prof. Tom Mens, chair of the ERCIM Working Group on Software Evolution. He started out with a presentation of the main research challenges in software evolution, based on feedback received from various members of the ERCIM WG.

Dr. Joost Visser represented the Software Improvement Group (SIG), a consultancy company based in Amsterdam, and presented an industrial perspective on software evolution challenges, based on informal feedback received from the many companies for which SIG provides software monitoring and risk assessment services. Prof. Ralf Lammel, from University of Koblenz in Germany talked about his former industrial experience while working at Microsoft, and the typical problems with software evolution he was confronted with in this industrial setting. Prof. Andreas Winter, head of the Software Engineering Group of the University of Oldenburg, Germany, stressed the importance of research in reverse engineering and migration of legacy systems, and explained that this research is bound to become more and more important, as technology is changing ever more rapidly, thus continuously giving rise to a new breed of legacy systems. Jim Buckley presented the vision of LERO, the Irish Software Engineering Research Centre. He stressed the importance of dynamic software adaptation, as programming languages are becoming more and more dynamic, and there is an industrial need of accommodating changes "on the fly".

These presentations allowed the European officers to get a more or less global view of the important issues that software engineers are continuously confronted with. As a result of the discussion that ensued after the presentations, all participants agreed that:

- evolving software is considerably more difficult than developing software from scratch, because the tools, processes, languages and mechanisms that are currently available are not sufficiently advanced
- one cannot properly deal with software evolution without taking a holistic view on the software system and everything that surrounds it
- the rapid changes in technologies and adoption of new paradigms only makes the software evolution problems become worse
- software evolution is not just a technical problem, it also concerns project managers and end-users.

Based on these insights, we believe that the main challenges and needs in software evolution are, in arbitrary order:

- better techniques (ie more automated and cost-effective) for software renovation, re-engineering, and migration of legacy systems
- better mechanisms to assess, control, improve and assure software product and process quality
- more sophisticated formalisms, tools and techniques to deal with software frameworks, product families and software product lines
- better techniques to support runtime evolution and dynamic software adaptation
- better managerial support to estimate and predict changes, effort and productivity
- automated support for software evolution and maintenance that is seamlessly integrated in software development environments
- solutions that scale up to very large distributed systems involving multiple languages and levels of abstraction, a large number of developers and users, and huge amounts of data
- more empirical and industrial validation of research results, as well a better industrial adoption of these results
- support for co-evolution, traceabililty and synchronisation between all artefacts used and produced during software development (including, but not limited to: code, data, tests, documentation, models, languages, specifications.)

To conclude, software evolution and maintenance is a major challenge for anyone confronted with software development, and the challenge is there to stay. If we want to address all the technological, managerial and social obstacles in this domain, the European research community in this field needs to join forces with European industry. Given the strong presence of European researchers in the field of software evolution and maintenance, and the importance for companies to excel in producing high quality software, we hope that these research challenges will be addressed.

Please contact:
Tom Mens, University of Mons
General Chair of ERCIM WG Software Evolution
E-mail: tom.mens@ercim.eu

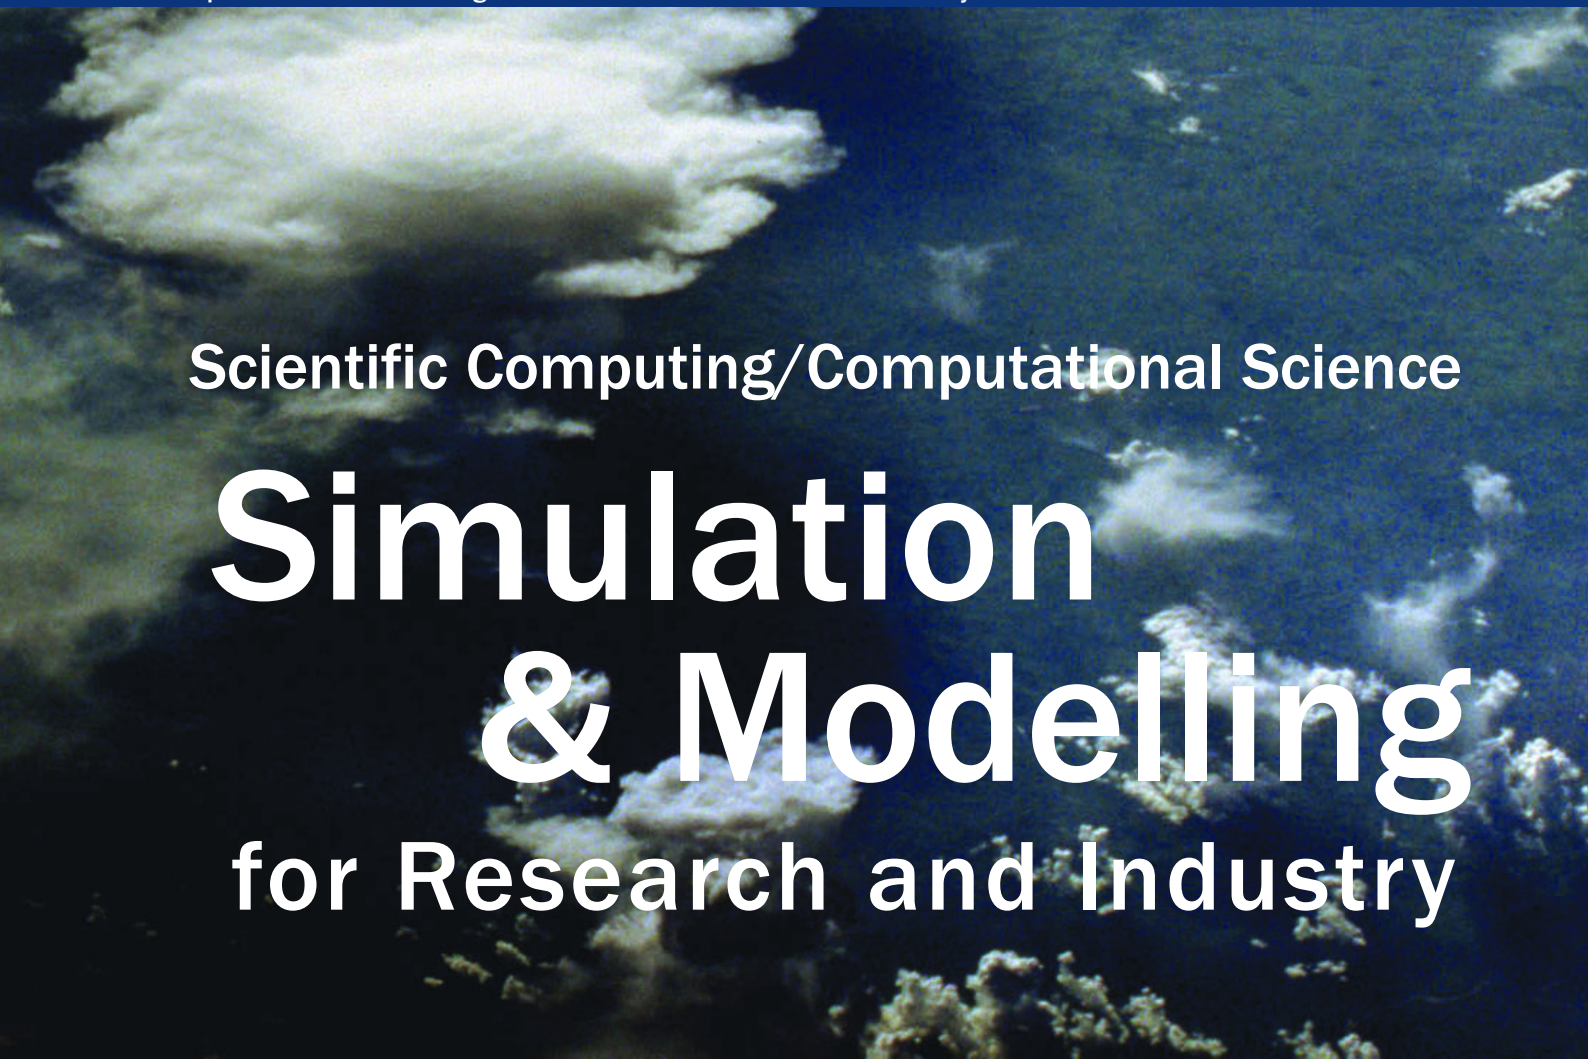

Scientific Computing/Computational Science

Simulation & Modelling

for Research and Industry

Introduction to the Special Theme

by Ulrich Trottenberg and Han La Poutré

Many phenomena and processes in nature, science, technology and economy are today modeled mathematically, and these models are used for control, prediction and optimization. Virtual models substitute real systems and simulation replaces costly, long lasting and dangerous experiments. Today problems can be treated on computers that seemed utopian to be solved twenty years ago.

Mathematics and Computer Science play a fundamental role in forming these simulated realities. In some sense the mathematical models are horizontally arranged in a landscape of vertically oriented disciplines and applications. Computer Science techniques then play a similar role, using or extending such models.
The articles in this special theme all deal with scientific and high performance computing, with computational science, with numerical and discrete simulation, and with large-scale computation and software – and they reflect a large variety of applications. The range of applications is indeed highly impressive.

The articles do not so much reflect the progress that has been and is being made by the hardware and computer development. We can all observe this progress in our everyday work and life, indeed. Much less, it is known that an even more dramatic progress has been made in mathematical modelling and algorithmic developments. An example: The computing times for a typical large-scale problem in scientific computing (a convection-diffusion type partial differential equation, discretisized in a grid with approx. one million grid points) can today be solved a million times faster than twenty years ago. Here a factor of less than 1000 is due to the hardware development and a factor of more than 1000 is due to the algorithmic progress. And this will continue. We need new algorithms and new algorithmic ideas in order to solve, for example, the global weather forecast problem. In 2012, the weather models, then discretized on a grid with 400 million grid points (to be solved several thousand times every day), are known and used and the 100 teraflops computers will be available – but the algorithms are missing. Therefore – much has to be done within the next two years.

Within Computational Science, the area of scientific computing is the oldest. This typically deals with physical problems, to be tackled with large-scale simulation. More recently, other scientific areas have arisen, from social and economic sciences, to the design of (computer) systems themselves, like in the engineering disciplines. Especially, large scale agent-based simulations systems more and more allow for more recent developments in the field: the simulation of socio-eco-

nomics systems. Here prediction and emergence play an important role. Similarly, intelligent techniques can deal with more complex optimization problems, like in supply chains or health care, due to large scale computation. In the engineering discplines, large-scale simulation is used to investigate eg the behaviour of the designed systems, like software (agent) systems.

At the same time, additional techniques from computer science are developed in order to deal with computational platforms, like advance visualisation techniques and data rendering and processing techniques.

And much could be done on the software side. The articles of the special theme show how much is being done and achieved in many fields of numerical simulation and computer science. Still – the output with respect to commercial and marketable software is small. Most of the developments remain prototypes; very few of these developments have the ambition to become real and successful software products for technical computing.

In this special issue, a wide variety of articles is included that give an impression of the area of computational science. This area goes from mathematics and scientific computing to advanced computer science techniques, like infrastructures, agent-based simulation, and visualisation. We hope the reader will get an impression of today's state of Computational Science / Scientific Computing (CS/SC) as multidisciplinary approaches in which advanced computing capabilities are used to understand and solve a wide variety of complex problems. As a methodology, CS/SC can rapidly be seen as the "third way" in science and engineering, allowing researchers to address questions that are largely inaccessible to experimental or theoretical investigations, and allowing practitioners to solve complex engineering and optimization problems.

For a strategic view on computational science, we like to recommend the PITAC 2005 report (PITAC: President's Information Technology Advisory Committee, USA): Computational Science: Ensuring America's Competitiveness. The report is available for download at: http://www.nitrd.gov/Pitac/reports/20050609_computational/computational.pdf

Please contact:
Ulrich Trottenberg
Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Germany
E-mail: ulrich.trottenberg@scai.fraunhofer.de

Han La Poutré
CWI, The Netherlands
E-mail: Han.La.Poutre@cwi.nl

The Forecast of Predictability and Instability in Physical Modelling

by Juan C. Vallejo and Miguel A. F. Sanjuán

All numerical calculations have inherent inaccuracies, and beyond certain timescales even the best method will diverge from the true orbit. The concept of shadowing time allows the reliability of a computer-generated orbit to be characterized. This indicator is being applied to some models of galactic potentials, where areas of high and low predictability mix in a fractal-like way.

In the last century, the numerical approach to solving physical problems has gained in relevance with the increase of computational facilities. Methods derived from chaos theory and nonlinear dynamics techniques are quite useful in solving real problems where chaos is present and a strong dependence on initial conditions is a key issue.

Predictability refers to the assessment of the likely errors in a forecast, either qualitatively or quantitatively. It refers directly to the stability/instability of the true orbit, but also to the coincidence of the calculated orbit itself, or reliability.

Instability is a very well-known factor when considering a given solution to the model. In a broad sense, an orbit is unstable if it is strongly dependent on initial conditions. More precisely, an orbit is chaotic if the dynamics is bounded, has at least one positive Lyapunov exponent and the omega set is not periodic and does not consist solely of equilibrium points and connecting arcs. The larger the instability, the larger is the likehood that the real orbit will diverge from the calculated one.

Typically, the reliability time goes with the inverse of the Lyapunov exponent, or Lyapunov time. As all numerical calculations have inherent inaccuracies, beyond certain timescales even the best method will diverge from the true orbit. The usual steps for analysing a given system involve calculating the instability for a given set of initial conditions. It happens that the standard definition of the Lyapunov exponent has a very long convergence time (if any). Due to the slow convergence of the asymptotic value, many others numerical indexes are used, but the basic idea remains the same.

However, we can also refer to the numerical reliability of the system understood as the confidence of the computed orbit and the true orbit during an interval of time, independently of its instability. Essentially, computers move from one true orbit to another nearby orbit because of rounding errors and internal floating number representations. However, the computed orbit (called pseudo-orbit) may still lead to correct predictions because of the existence of a nearby exact solution. Otherwise, the calculated orbit may be very distant.

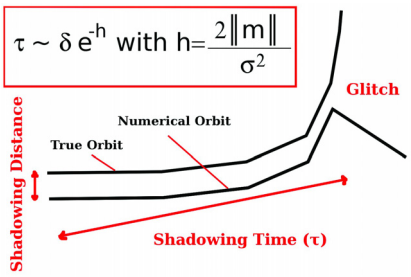

Figure 1: Computers move from one true orbit to another nearby orbit because rounding errors and internal floating number representations. How long the true orbit is shadowed by the computed orbit before suffering a glitch and later deviation can be estimated from Lyapunov exponent distributions.

The true orbit is called a shadow, and the existence of shadow orbits is a very strong property, with the shadowing time being a valid limit for the predictability of the system. A basic requirement for shadowing is that the system be hyperbolic. In case of nonhyperbolicity, the point may not be shadowed and the computed orbit behaviour may be completely different from the true one.

Stability is not the same as hyperbolicity. An orbit can be unstable and yet hyperbolic. When the orbit is nonhyperbolic, there can be still a moderate obstacle to shadowing when the nonhyperbolicity is due to tangencies between stable and unstable manifolds. But in the so-called pseudo-deterministic systems, when nonhyperbolicity is sourced to a phenomenon named unstable dimension variability, the shadowing is

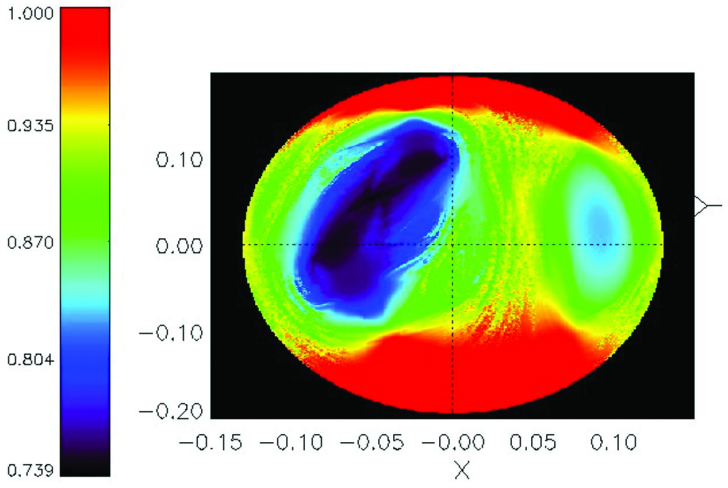

Figure 2: Plotting the instability and hyperbolicity indexes provides insight into the predictability for every initial condition of the available phase space. However, such indexes distribute in a very complex way under certain conditions of the model parameters, often showing fractal structures like those in the figure.

only valid during trajectories of given lengths, which may be very short.

Our work in the Nonlinear Dynamics, Chaos and Complex Systems Group at the Universidad Rey Juan Carlos (URJC) in Madrid, is focused on analysing and plotting predictability charts for certain models. We get the shadowing times by computing the finite time Lyapunov exponent distributions. Each computational run is based on the integration of one initial condition with a well-known Runge-Kutta-Fehlberg integrator scheme. As the integration progresses, the distribution is calculated and the instability and hyperbolicity indexes are derived. This is repeated for several conditions and parameters, thus returning a complete description about the predictability of the used model.

Note that this strategy is very well suited for high-throughput (Grid) numerical schemas. A set of shell scripts feeds the integrator with the proper initial conditions and a grid engine (in our case, GridWay) and submits the process to the proper cluster in a user transparent way.

These techniques are general enough to be applied to different models with minor modifications. Our work focuses on galactic modelling, where the applied timescales are critical. Mathematically meaningful timescales cannot be physically acceptable here. A particle being shadowed during only a few crossing times (cycles) is a critical issue. Galaxies are evolutionary entities, from the point of view of both their gravitational potential and their constituents. For instance, in a Milky Way type model, the older stars will have orbited the galactic centre just sixty times before death. And the galaxy itself may evolve during this short period of time. Interpreting the dynamical system in a broad sense, given an initial condition, the stability refers to the location itself, not to the tracer particles. But considering what happens to the model in such short timescales, and for the sake of our discussion, the forecast of the numerical galactic model reliability is then of importance.

Link: http://www.fisica.escet.es

Please contact:
Juan C. Vallejo
Universidad Rey Juan Carlos, Madrid, Spain
Tel: +34 91 8131366
E-mail: juancarlos.vallejo@urjc.es

Realistic Material Appearance Modelling

by Michal Haindl, Jiří Filip and Martin Hatka

Physically correct and realistic visual appearance rendering or analysis of material surface visual properties require complex descriptive models capable of modelling material dependence on variable illumination and viewing conditions. While recent advances in computer hardware and virtual modelling are finally allowing the view and illumination dependencies of natural surface materials to be taken into account, this occurs at the expense of an immense increase in the required number of material sample measurements. The introduction of fast compression, modelling and rendering methods for visual data measurements is therefore inevitable.

The established practice in computer vision, graphics and pattern recognition is to base inferences only on restricted information. Virtual reality applications typically use oversimplified models that cannot even remotely approximate the appearance of real scenes, meaning human observers can easily differentiate between real and virtual scenes. Fortunately, the recent swift development of computer technology has allowed models and tools that seemed theoretical only a decade ago to become feasible as a part of foreseeable future routine processes. Physically correct and accurate material appearance visualization is in high demand. It not only has a huge economic impact in visual safety simulations and virtual design in automotive industry and architecture, but also has large potential in visual scene analytical applications, including health care, security, defect detection and content-based image retrieval. However, many challenging problems still exist, such as efficient measurement of material optical properties, image compression, optimal mathematical representation, unsupervised segmentation and interpretation and many others.

We have developed several multidimensional probabilistic models based either on a set of underlying Markov random fields or probabilistic mixtures, which allow physically correct surface lossless representation and modelling, huge measurement space compression (so far unbeaten at up to 1:1 000 000), and even modelling of previously unseen surface data or their editing. These methods are parametric, so they do not require original measurements to be stored. However, such models are nontrivial and suffer from several challenging theoretical problems such as stability, parameter estimation and non-iterative synthesis, which must be circumvented.

Alternative approaches using physical reflectance models or sophisticated sampling were also investigated. Regardless of the traits of individual models, they all meet comprehensible requirements such as unlimited seamless surface image enlargement, high visual quality and compression, as well as some less obvious features like strict separation of the analytical and synthesis parts, possible parallelization and implementation in advanced graphics hardware. Unfortunately, there is no ideal universal visual surface mathematical model suitable for all applications or material types. Each of these aforementioned models have their advantages and drawbacks simultaneously, hence optimal measurement as well as modelling depends on both material and intended application, and must be automatically recognized. Surprisingly, the reliable assessment of visual quality is also a difficult task because no usable mathematical crite-

rion exists and such verification requires costly and time-demanding psychophysical visual evaluation. On the other hand, we have successfully applied methods of visual psychophysics to the development of even more efficient material-dependent compression and measurement methods.

We believe that the combination of perceptually optimized measurement and effective mathematical representation of surface appearance is a key to the wide applied usage of realistic view- and illumination-dependent surface material optical measurements.

Links:
http://ro.utia.cas.cz/10cvpr_tutorial/
http://staff.utia.cas.cz/filip/projects/pertex

Please contact:
Michal Haindl - CRCIM (UTIA)
Tel: +420 266052350
E-mail: haindl@utia.cas.cz

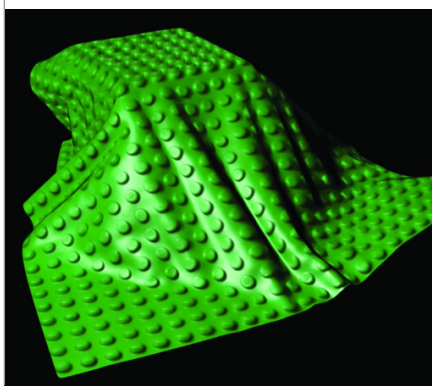

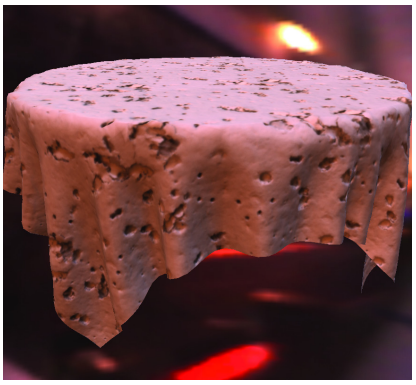

Figure 1: Examples of the realistic rendering of view- and illumination-dependent textures in a car interior, on virtual drapery and on an environmentally lit tablecloth.

Risk Assessment Study based of the 365AD Earthquake to Drive a Civil Protection Exercise

by Evangelia T. Flouri, Costas E. Synolakis, Nikos A. Kampanis and Catherine E. Chronaki

The simulation of the extreme destructive historical earthquake of 365 AD in the Aegean Sea will help to draw a realistic disaster scenario to test EU member state collaboration in the frame of the European Civil Protection Mechanism.

POSEIDON, a project co-funded by the European Commission, DG Environment, Civil Protection Unit, aims to design, organize and evaluate an operational civil protection exercise involving Greece, Cyprus and France, addressing the results of a hypothetical extreme disaster caused by an earthquake of considerable magnitude followed by a tsunami. This hypothetical disaster will devastate the island of Crete and will require the invocation of the European Civil Protection Mechanism to bring assistance from other areas in Greece, France and Cyprus. The aims of POSEIDON are: (a) to understand, verify and improve civil protection procedures for preparedness and response, particularly in relation to earthquake events followed by tsunamis, (b) to provide a learning opportunity for the actors involved in civil protection assistance interventions, (c) to test the coordination of EU civil protection assistance particularly through the Monitoring and Information Centre (MIC), and (d) to validate new applications for early warning and communications systems as well as procedures by which to inform the public of emergency measures to be undertaken.

FORTH-IACM, and specifically the Laboratory of Coastal Research & Applications is responsible for performing a risk assessment study that will provide the hypothetical disaster scenario from a large earthquake and the tsunami that follows, using advanced simulation techniques. The risk assessment study involves the production of tsunami inundation and seismic intensity maps to be used before the exercise by the emergency planners to prepare their response and then during the operational exercise when additional local scenarios will be added to overwhelm the local command and control and direct search and rescue operations.

We take as our starting point the historical earthquake of 365 AD that occurred in the sea west of the island of Crete and which was followed by a tsunami. Historical sources report that this particular earthquake caused extensive destruction, and this has been verified by contemporary seismological and geological research (Shaw B, Ambraseys NN, England PC, Eastern Mediterranean tectonics and tsunami hazard inferred from the AD 365 earthquake, Nature Geoscience 1, 268 - 276 (2008)). Apart from the destruction, the earthquake caused significant geological disturbances, particularly acute in the westward coast of the island.

The reconstruction of the earthquake characteristics from bibliographical

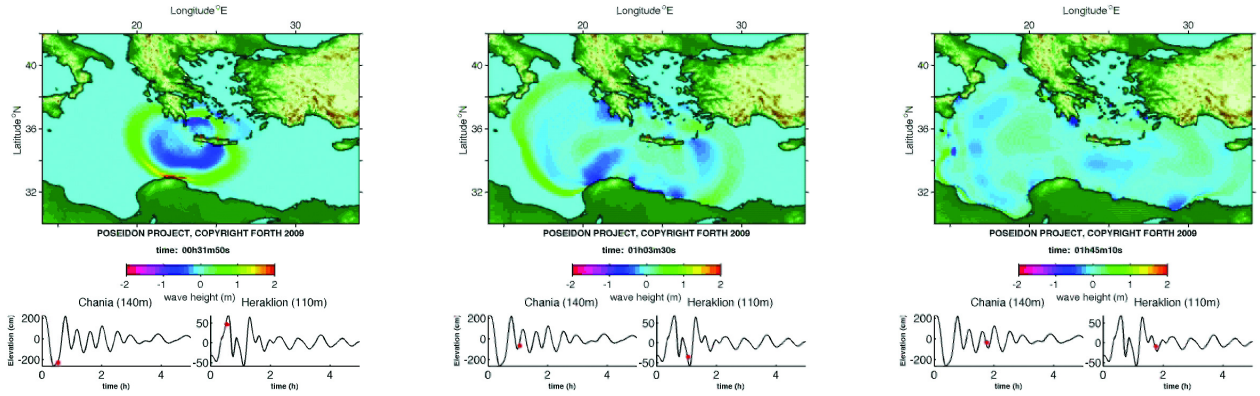

Figure 1: Snapshots of the propagation of the tsunami generated by the 365 AD earthquake.

information will simulate the movement of the tectonic plates in the rupture area, and will provide the initial conditions for the numerical simulation of the tsunami wave. The software used performs a full simulation of the tsunami wave, that is, its generation, propagation in the open sea, and run-up when it reaches the coast (see article by Evangelia Flouri et al on page 16). The exact bathymetric data around Crete were acquired from the Hydrographic Service of the Hellenic Navy, for the purposes of more realistic modelling.

FORTH-ICS will be responsible for the deployment and evaluation of information and telecommunication technology in the POSEIDON exercise, emphasizing the emergency deployment of auxiliary satellite and WiFi networks, to support the strained telecommunication infrastructure. Relevant standards such as those proposed by OASIS, W3C and HL7 will also be reviewed and evaluated in the context of the exercise, focusing on the timely provision of information to the public as well as on efficient and effective management of emergency episodes.

Interoperability considerations are key to the POSEIDON project, since one of its main aims is to demonstrate effective cross-border collaboration between France, Greece and Cyprus, possibly identifying issues related to the use of the civil protection modules in the context of the European Civil Protection Mechanism.

FORTH-ICS has been one of the key participants along with Institut de Médecine et de Physiologie Spatiales - MEDES (FR) and Telemedicine Technologies (FR) in Satellites for Epidemiology SAFE, a 2007 civil protection exercise, aiming to demonstrate the value of satellite-enabled applications for early warning in public health following an epidemic outbreak, with the backdrop of a severe earthquake (see article in ERCIM 75). With POSEIDON we move a step forward, as scientific risk assessment based on a historic earthquake disaster is used to create a realistic disaster scenario to test EU member state collaboration in the frame of the European Civil Protection Mechanism.

The Project is coordinated by Mrs. Koutentaki, Director of Civil Protection in the Region of Crete, Greece. The Civil Defence of Cyprus (Cy), the General Secretariat of Civil Protection (GR), the Hellenic Red Cross (GR), Civil Protection Sans Frontiers (FR), Telemedicine Technologies (FR) and FORTH (GR) are partners in the project. FORTH participates in the POSEIDON project with two of its seven institutes, the Institute of Computer Science (FORTH-ICS) and the Institute of Computational Mathematics (FORTH-IACM).

Civil protection exercise.

Links:

Region of Crete, Directorate of Civil Protection: http://www.crete-region.gr

FORTH-Institute of Computer Science: http://www.ics.forth.gr

FORTH-Institute of Computational Mathematics: http://www.iacm.forth.gr

General Secreteriat of Civil protection, Ministry of Citizen Protection, Greece, http://www.gsrt.gr

Civil Defence, Ministry of Interior, Cyprus: http://www.moi.gov.cy/moi/cd/cd.nsf/intro?OpenForm

European Civil Protection: http://ec.europa.eu/environment/civil/

Hellenic Red Cross: http://www.redcross.gr/

FORTH-IACM, Laboratory of Coastal Research & Applications: http://www.iacm.forth.gr/CFD/CIRU.php

Please contact:
Nikos Kampanis
FORTH-IACM, Greece
E-mail: kampanis@iacm.forth.gr

Catherine Chronaki
FORTH-ICS, Greece
E-mail: chronaki@ics.forth.gr

Mathematical Modelling of Tsunami Waves

by Evangelia Flouri, Dimitrios Mitsoudis, Nektarios Chrysoulakis, Manolis Diamandakis, Vassilios A. Dougalis and Costas E. Synolakis

Tsunami waves (long waves) can be efficiently simulated by numerical models solving the Shallow Water equations. Recently, FORTH-IACM has used depth-averaged computational models in shallow water with an emphasis on complex 3D domains in the context of some EU projects. One of these was TRANSFER (Tsunami Risk ANd Strategies For the European Region).

The main goal of the TRANSFER project - an FP6 project, with 30 partners from 12 countries - was to contribute to the understanding of tsunami processes in the Euro-Mediterranean region, to improve numerical models, the vulnerability and risk assessment from tsunami hazards, and to identify best strategies for reducing tsunami risk. For the application of methodologies, seven test areas in different countries were selected as test cases, from among which the area of Rhodes-Fethiye in the south-east Aegean Sea was chosen as the 'master test case' of the project. For all these areas, innovative probabilistic and statistical approaches for tsunami hazard assessment were applied as well as up-to-date techniques for computing inundation maps. Many tsunami scenarios were envisaged for these areas; vulnerability and risk were assessed, and prevention and mitigation measures were defined.

FORTH-IACM participated in TRANSFER with significant contributions in various tasks, and was the leader (responsible institution) for the production of tsunami inundation maps. All methodologies developed by FORTH-IACM were applied to the master test case, focusing on the Greek island of Rhodes, and especially on the capital city (Rhodes). Under this task, the work of FORTH-IACM was the production of inundation maps along the coast of the city for large tsunami scenarios, the production of a detailed vulnerability and risk analysis and the development of tsunami impact prevention and mitigation strategies.

The tsunami hazard estimation was based on numerical modelling using a full tsunami simulation (generation, propagation, and inundation) with the MOST (Method of Splitting Tsunamis) numerical model (Titov & Synolakis, 1998), coupled with updated and highly accurate bathymetry and topography data. The model implements an elastic fault plane mode to calculate an initial condition, which in turn is propagated across the deep sea by solving numerically the Nonlinear Shallow Water equations with a split-step, characteristics-finite difference scheme. Finally, the coastal flooding and inundation is simulated by extending the propagation calculations with the aid of a multi-grid

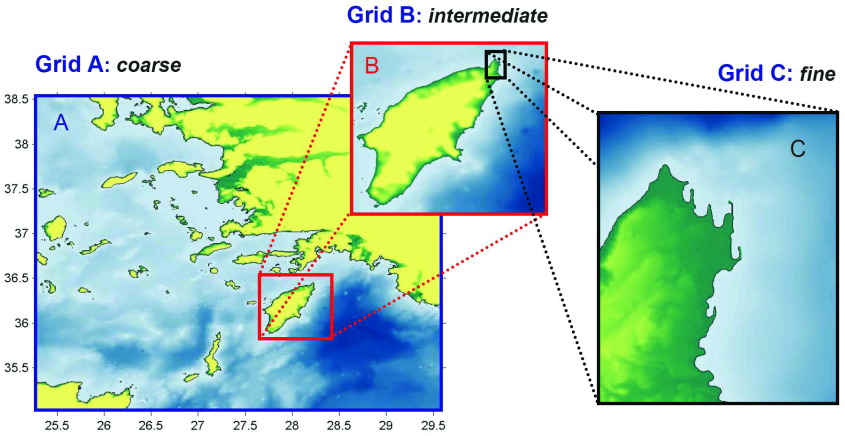

Figure 1: The study area and the areas covered by the three nested grids used by the numerical model.

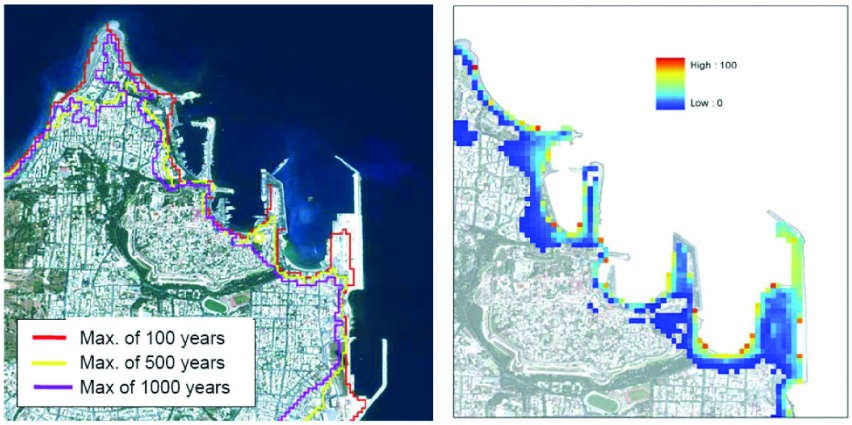

Figure 2: Probabilistic maps for the city of Rhodes: inundation lines for 100, 500 and 1000 years superimposed on an IKONOS satellite image (left); frequency of inundation among 100 runs for 1000 years depicted in colours over the satellite image (right).

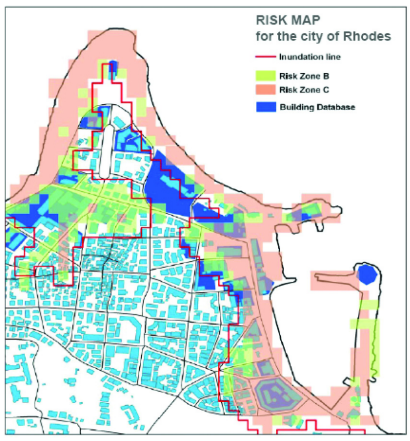

Figure 3: Tsunami risk map for the city of Rhodes (northern part of the port), with the different risk zones according to the flow depth and the inundation line.

run-up algorithm. The code requires as input detailed information on seismic source mechanisms, gridded bathymetric data information for the open sea propagation, and a set of gridded Digital Elevation Models (DEM) containing bathymetry and topography for use during the inundation phase. Accurate bathymetry and topography data for Rhodes were purchased/collected from various sources, GIS-processed and coupled in the nested bathymetry/topography grids shown in Figure 1.

Hundreds of simulations were performed based on both a deterministic and a probabilistic approach and various types of tsunami inundation (inland penetration) maps were produced (see Figure 2). The deterministic maps have been produced from simulations based on worst-case scenarios with the aim of assessing the impact of unusually large (hypothetical) seismic events. The probabilistic maps were the outcome of simulations based on multiple near-field tsunami scenarios and took into account the uncertainty associated mainly with the location of the seismic source. The comparison of the various deterministic and probabilistic scenarios led to the conclusion that the city of Rhodes is most heavily affected by the probabilistic scenario of a 1000-year event. All conclusions and results were utilized for the needs of the vulnerability and risk analysis.

For the purposes of the latter task specific basic data were identified, collected and analysed in order to obtain estimates of vulnerability and risk. To this end, elements such as land cover/use, buildings and coastal structures and infrastructure (roads, ports, piers) were identified, as well as tourist populations with high seasonal variation. All these relevant data were combined with the results of simulations and used to produce flooding risk maps and their representation within a GIS. The estimation of tsunami risk was based on the inundation maps as well as on the computed values of the flow depth in the coastal areas. Tsunami risk zones were defined and presented in GIS layers and maps, see Figure 3. Based on this analysis, the FORTH-IACM team developed and proposed detailed recommendations.

Link:
http://www.transferproject.eu/

Please contact:
Evangelia Flouri
IACM-FORTH, Greece
E-mail: flouri@iacm.forth.gr

Comprehensive Modelling and Simulation System for Decision Support in the Field of Radiation Protection

by Petr Pecha and Radek Hofman

Potential failures in man–made processes can result in the accidental release of harmful substances into the environment. Risk evaluation and a decision-making process that is focused on protecting the population has the highest priority. Historically, accidents in nuclear facilities have revealed a lack of sufficiently advanced decision support software tools. Great attention has been paid to this topic since the Chernobyl disaster. The software system HARP (HAzardous Radioactivity Propagation) is designed for the fast assessment of the radiological consequences of such a release of radionuclides into the environment.

The HARP system is the application part of a grant project supported by the Grant Agency of the Czech Republic (period 2007–2009), which was solved in the Institute of Information Theory and Automation of the Academy of Sciences of the Czech Republic. The new version of the product is a complex software tool for modelling the radiological consequences of radionuclide releases due to the normal and emergency operation of a nuclear facility. Aerial transport of discharged radionuclides is studied up to a 100km radius from the pollution source. Dispersion, deposition and successive radioactivity migration towards the human body is simultaneously modelled. As the system contains a database of demographical data on the area of interest, it can evaluate the major radiological quantities and the radiological burden on the population due to different pathways of irradiation in both the early and the late phases of a radiation accident.

The core of the system is an atmospheric dispersion model. Generally, the modular architecture of the system enables an arbitrary dispersion model to be inserted. The default is the segmented Gaussian plume model. Although simple, the Gaussian model is consistent with the random nature of atmospheric turbulence and is an approximate solution of the Fickian diffusion equation. Proven semi-empirical formulas are available for the approximation of important effects like the interaction of the plume with proximal buildings, momentum and buoyant plume rise during release, power-law formula for estimation of wind speed changes with height, depletion of the plume activity due to the removal processes of dry and wet depositions, radioactive decay and creation of daughter products, dependency on physical-chemical forms of admixtures and land-use characteristics, plume lofting above the inversion layer, and so forth.

A special emphasis is laid on the proper treatment of types of input parameter fluctuations, in the sense of differentiation between variability and uncertainty. Some model uncertainties arising from the conceptual limitations can be roughly estimated on the basis of an

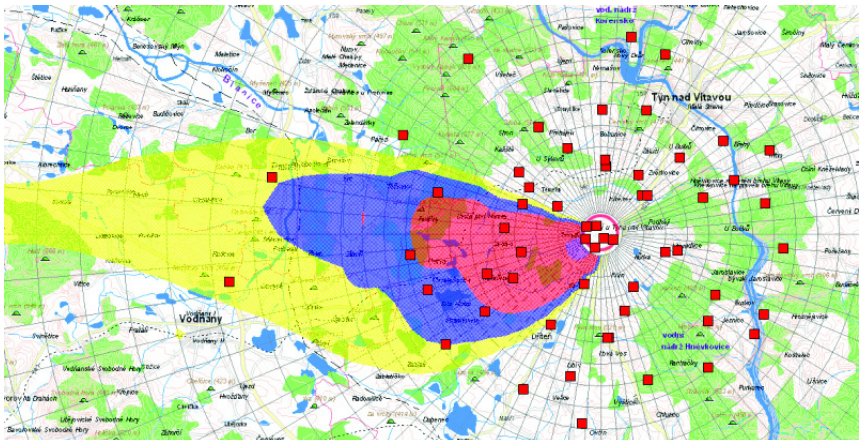

Figure 1: Prediction of terrain contamination based on meteorological forecast due to a continuous release of mixture of radionuclides just after the three hours since the release start. Red squares denote locations of fictitious measuring stations providing measurements for assimilation with the model predictions.

ensemble-based approach with alternative parameterization of physical effects taking place in the atmospheric dispersion. The system offers an extensive interactive graphical user interface for presenting a wide range of outputs important for decision makers. Thus, the system also comprises a simulation and training tool enabling responsible staff to improve their knowledge and perception of the problem details. The HARP code has proved useful in the fields of evaluation of environmental impact assessment (EIA) and probability safety assessment (PSA) studies, where the influence of operation of a nuclear facility on the surrounding environment is appraised.

Advanced data assimilation methods based on Bayesian filtering developed within the grant project are incorporated into the assimilation subsystem. This means that the system offers a framework for the embodiment of relevant information from different sources, such as measurements and expert judgements, in a statistically optimal way. Provided that the system is connected to a monitoring network, it can be run in the online regime, where the subjectively chosen parameters regarding the release scenario (magnitude of release etc) are iteratively re-estimated upon measurements. This assimilation methodology can be also used as a tool to test the different topologies of a monitoring network and select the best one with regard to its functionality and to economical and other constraints.

The HARP system is tuned and tested in cooperation with National Radiation Protection Institute of the Czech Republic in Prague, where the product is connected to a database server providing up-to-date short-term meteorological forecasts on a three-dimensional grid. Exploitation of detailed meteorological data further improves the reliability of the embedded dispersion model.

Link:
http://havarrp.utia.cas.cz/eng

Please contact:
Petr Pecha
CRCIM (UTIA), Czech Republic
Tel: +420 266052009
E-mail: pecha@utia.cas.cz

Massively Parallel Simulations of Star-Forming Gas Clouds

by Stephen O'Sullivan and Turlough Downes

Many of the stars in our universe form inside vast clouds of magnetized gas known as plasma. The complexity of these clouds is such that astrophysicists wishing to run simulations could comfortably use hundreds of thousands of processors on the most powerful supercomputers. In the past, a serious obstacle to capitalizing on such computational power has been that the methods available for solving the necessary equations were poorly suited to implementation on massively parallel supercomputers.

Background: The Physics of Multifluid Gases

Usually, to study how a plasma behaves when a magnetic field is present, astrophysicists assume a single fluid with the field firmly anchored into it: if the plasma moves then so does the field, and vice versa. This picture makes some intuitive sense since charged particles try to travel along magnetic field lines – this is the principle used in older cathode-ray tube TV sets to direct the electrons onto the phosphor inside the screen.

The single-fluid picture, however, is only approximate. Most real plasmas consist of different families of particles, each of a different size and charge. These families will behave differently when interacting with the same magnetic field and with the rest of the gas. Indeed, many particles in astrophysical plasmas have no charge and so feel no direct effect of the magnetic field at all!

So, in many environments the physics is far richer than the simple single-fluid approximation: electrons, ions, neutral particles and even electrically charged dust grains can move differently. The plasma can no longer be said to act as a single coherent fluid with a tied-in magnetic field. As a consequence, the field may spread and twist in complex ways in response to the differing flows of the different families of particles. Under these circumstances, it becomes necessary to adopt a true multifluid picture of the system.

It is the necessity of adhering to the multifluid model which has made the prospect of running simulations of star-forming gas clouds on large-scale supercomputers prohibitively difficult in the past.

The HYDRA Project

Since 2004, Dr Stephen O'Sullivan and Dr Turlough Downes have been advancing a wholly new approach to the simulation of astrophysical gas clouds on parallel supercomputers. HYDRA, the principal code under development during that time, has evolved from a bleeding edge one-dimensional single-processor prototype to a three-dimensional production class code with successful runs completed on almost 300,000 cores of the Jülich supercomputer JUGENE.

The essence of the numerical approach employed by HYDRA is that, while the physics of multifluid interactions is complex, the resulting effects on fluctuations in the magnetic field can be classified into just two types: dissipation, which smooths the field; and dispersion, which separates fluctuations at different scales. Crucially for HYDRA, by appropriate treatment of the governing equations, these effects may be considered individually. The numerical challenges of integrating the equations are then found to be more manageable: a previously obscure technique known as Super Time Stepping has been adapted from the literature to treat dissipation while a completely novel scheme called the Hall Diffusion Scheme has been developed to deal with dispersion.

Figure 1 shows a recent run from a simulation which includes the full multifluid physics of turbulence in star-forming clouds. The simulation is carried in a cube made up of 512^3 computational zones. Such a simulation is very computationally demanding and can take up to twelve days to run using 4096 cores.

Until recently, most of the simulations run under this project were carried out for idealized scenarios with initially turbulent flows that were allowed simply to relax without the influence of any effects which might stir the gas up. Current efforts are focused on investigating the consequences of randomly mixing the flow. Under such circumstances, the kinetic energy added via mixing may ultimately end up increasing the strength of the magnetic fields. This so-called dynamo effect may have a profound influence on the evolution of star-forming clouds in nature. In addition, the mixing of the cloud can prevent or, perhaps counter-intuitively, actually enhance the prospect of stars like our own sun forming.

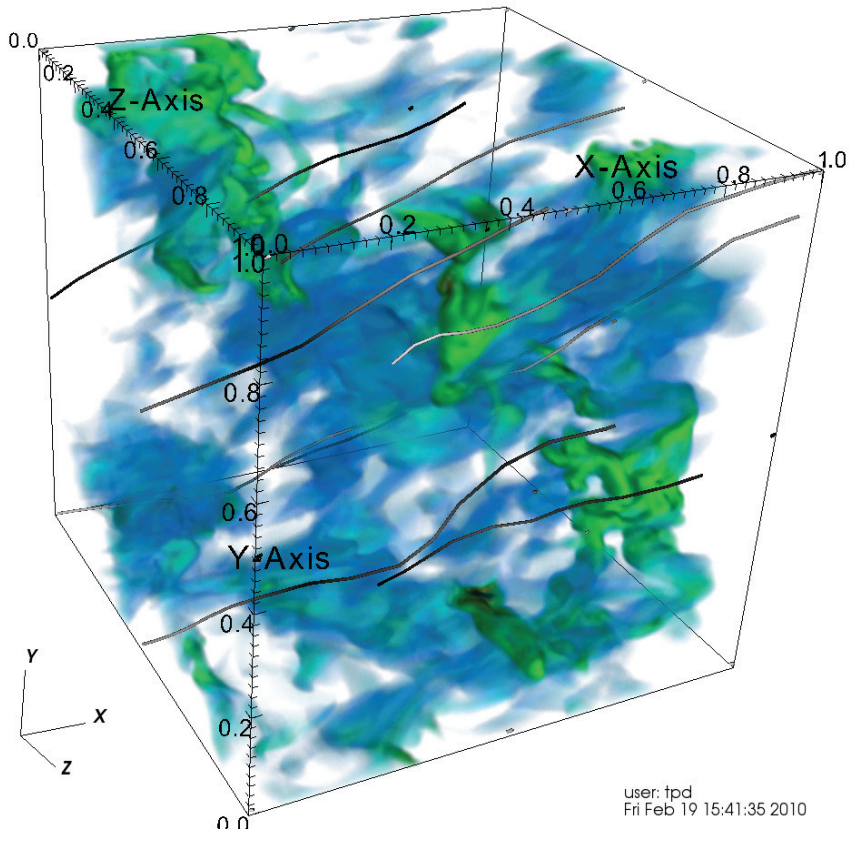

Figure 1: Snapshot from simulation of decaying turbulence in star-forming gas cloud.

At present, the principal institutions involved in this project are the Dublin Institute of Technology (School of Mathematical Sciences), Dublin City University (School of Mathematical Sciences & National Centre for Plasma Science and Technology) and the Dublin Institute for Advanced Studies (School of Cosmic Physics). Recent support has been provided by Science Foundation Ireland, under the Research Frontiers Programme, and the Partnership for Advanced Computing in Europe (PRACE). Support from the Higher Education Authority under Cycle 3 of the Programme for Research in Third-Level Institutions (PRTLI 3) was important during the early development phase of the code.

Central to the progress that has been made in this project are state-of-the-art Irish computational facilities, including the supercomputing 'capability' facilities of the Irish Centre for High-End Computing (ICHEC). Part of the code development was also carried out using facilities at the Argonne National Laboratory, USA.

Links:

http://www.dit.ie
http://www.dcu.ie
http://www.ncpst.ie
http://www.dias.ie
http://www.ichec.ie
http://www.prace-project.eu
http://www.cosmogrid.ie
http://www.sfi.ie
http://www.hea.ie
http://www.anl.gov

Please contact:
Stephen O'Sullivan
School of Mathematical Sciences,
Dublin Institute of Technology, Ireland
Tel: +353 1 402 4823
E-mail: stephen.osullivan@dit.ie

Turlough Downes
School of Mathematical Sciences,
Dublin City University, Ireland
Tel: +353 1 700 5270
E-mail: turlough.downes@dcu.ie

Scientific Computing for Astrophysical Map Analysis

by Emanuele Salerno

A research team at the Signal and Image Processing Lab of ISTI-CNR has been involved in studying data analysis algorithms for the European Space Agency's Planck Surveyor Satellite since 1999. The huge amount of data on the cosmic microwave background radiation provided by the Planck sensors requires very efficient analysis algorithms and high-performance computing facilities. The CNR group has proposed some of the source separation procedures that are now operational at the Planck data processing centre in Trieste, Italy.

Satellite-borne radiometric observations for cosmic microwave background studies are normally presented as sets of multichannel, all-sky maps where the background signal is contaminated by a number of galactic and extragalactic sources that are dominant in some frequency regions. When the measurements are particularly sensitive, the task of separating all the radiation components in the maps is particularly important. This is the case with the ESA cosmological mission known as the Planck Surveyor Satellite, whose nine channels, centred on frequencies between 30 and 857 GHz, provide huge all-sky maps with unequalled accuracy and angular resolution (see also ERCIM News 49 p. 14).

The Planck collaboration includes hundreds of scientists and dozens of institutes from all over the world. The satellite carries a telescope, provided by the Danish National Space Centre, with two instruments in its focal plane: a three-channel radiometer called the Low Frequency Instrument, under the responsibility of the Italian National Institute of Astrophysics in Bologna, and a six-channel bolometric sensor called the High Frequency Instrument, under the responsibility of the French Institute for Space Astrophysics in Orsay. The spacecraft was sent into orbit by an Ariane launcher on 14 May 2009, from the European spaceport in Kourou, French Guiana. The mission has now completed its first coverage of the microwave sky; the end of operations is expected in 2011.

Studying the small anisotropies of the cosmic background is a fundamental step in the understanding of many aspects of the origin and the evolution of the universe. As the source separation problem is of crucial importance for the scientific goals of Planck, it has been studied intensely by many groups. The Planck group at ISTI-CNR in Pisa has been working on it since 1999, when the project's first plenary meeting was held in Capri, Italy.

Our group started by adopting blind source separation techniques, motivated by the lack of an accurate data model and by a provisional assumption of mutual independence between the interfering radiations. In successive years, more specific techniques have been developed, taking into account pieces of information neglected by purely blind methods. In particular, the data model was enriched by possible dependencies between sources along with their intrinsic spatial structures. The non-stationarity of the mixing operator and the instrumental noise have also been considered. Finally, as this is always important in scientific data analysis, accurate methods for error estimation have been devised. The computational techniques adopted are diverse, from quick, fixed-

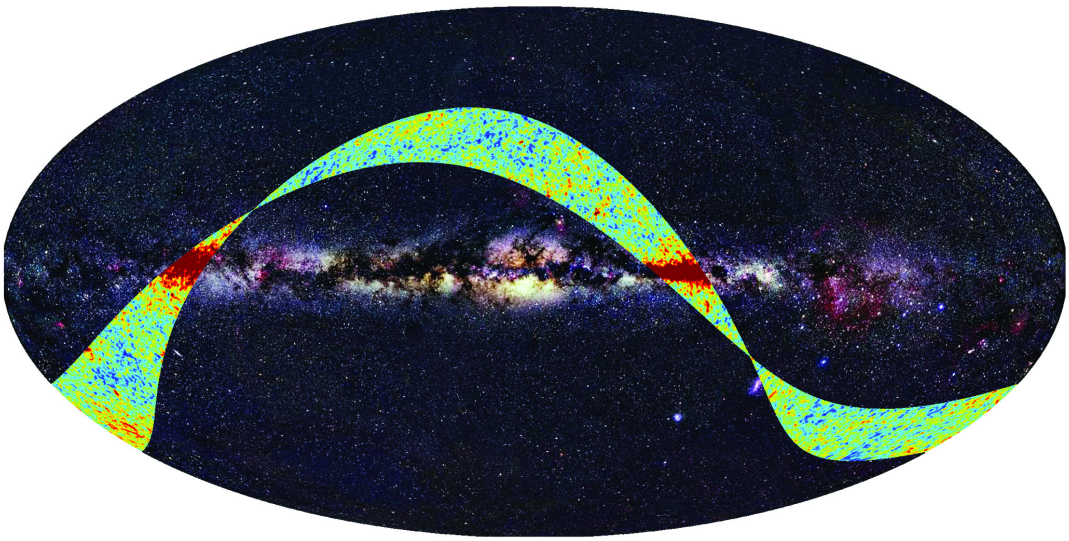

Figure 1: The first microwave sky data released by the Planck mission, superimposed to an all-sky map at optical wavelengths. The bright equatorial area represents the intense radiation from our Milky Way galaxy. The microwave data are mapped through a colour scale that indicates the deviations of the radiation temperature from the average background value of 2.726 K (red is hotter and blue is colder). Date: 17 Sep 2009 Satellite: Planck. Copyright: ESA, LFI & HFI Consortia (Planck), Background image: Axel Mellinger.

point optimization strategies to accurate but lengthy fully Bayesian approaches. Since a single Planck map can have around $5x10^7$ pixels, even very fast algorithms need supercomputing facilities to perform a task on the complete data set. This is why the Planck data-processing centres in Trieste and Paris are equipped with very fast and powerful hardware, and only the fastest component separation algorithms have been implemented in the routine data analysis pipelines.

Nevertheless, the more sophisticated Bayesian methods are still being studied, since they potentially offer advantages in terms of flexibility, capability of including prior knowledge, and accuracy in results and error estimates. Happily, the computational complexity of some of the new versions we are developing is decreasing, and we expect that further refinements will make them suitable to be applied to large amounts of data. Even if these algorithms are not ready to be used routinely before the end of the mission, their features can enable the research teams to perform accurate analyses on at least part of the complete data set.

The group in Pisa has been collaborating with several other Planck groups, mainly at the International School for Advanced Studies in Trieste, the Astronomical Observatory in Padova, Italy, and the Cantabria Institute of Physics in Santander, Spain. Some of the stochastic algorithms that we propose have been designed in collaboration with the ERCIM partners at Trinity College Dublin, in the framework of the ERCIM-led MUSCLE European Network of Excellence (see also ERCIM News 71 p. 6).

Links:
Main Planck portal:
http://www.esa.int/SPECIALS/Planck/index.html

Planck science homepage:
http://www.rssd.esa.int/index.php?project=PLANCK

Please contact:
Emanuele Salerno
ISTI-CNR
Tel: +39 050 315 3137
E-mail: emanuele.salerno@isti.cnr.it

Modelling the Growth of a Malignant Brain Tumour

by Alexandros Roniotis, Kostas Marias, Vangelis Sakkalis

One of the major aims of the ContraCancrum Project is to develop a composite multilevel platform for simulating glioma development as well as tumour and normal tissue response to therapeutic modalities and treatment schedules. By efficiently predicting the evolution of a tumour and how this alters with different therapeutic schemes, clinicians could optimize the disease treatment procedure in the patient's individualized context.

Glioma is the most malignant type of brain cancer. Glioma cells invade into neighbouring normal tissues so rapidly that all common imaging techniques (MRI, CT, PET) cannot detect them. Even if a tumour is detected using medical imaging, there will be glioma cells dispersed beyond this area. Thus, if the clinicians apply therapeutic schemes only in the imaged tumour, the remaining cells around it will still proliferate and cancer will recur. Moreover, leaving safety margins around the imaged bulk is inefficient, because glioma migration is anisotropic, being facilitated along white fibres. Apart from this, the speed of migration in white matter is five times faster than in grey matter.

The ContraCancrum Project develops diffusive models to simulate the diffusive behavior of glioma, taking anisotropy and matter heterogeneity of invasion into account. These models use medical data (DICOM images) taken

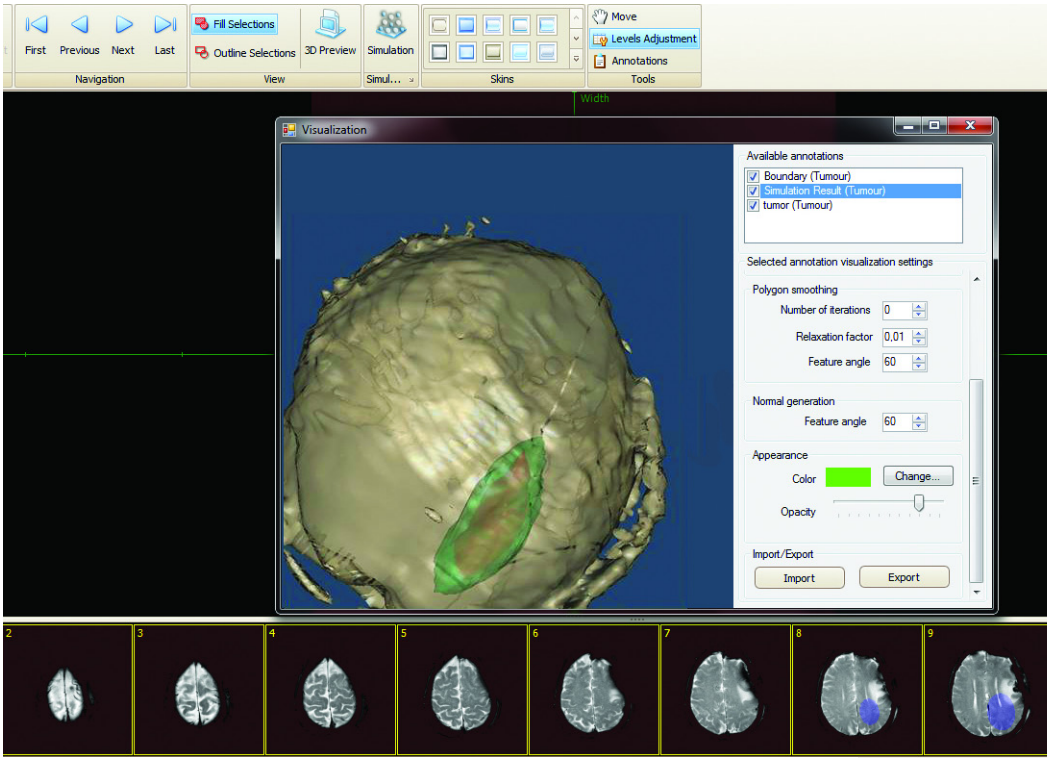

Figure 1: An example of visualizing the simulated tumor in brain, using the DoctorEye Tool, after 65 days of evolution. The initial tumor is reddish, while the simulated tumor is greenish.

from the Oncologic Clinic of Saarland University in Germany. Tumour areas and initial glioma concentration are initially annotated by radiologists, using the DoctorEye application. DoctorEye is an open-source platform developed by FORTH, and is being used for fast and precise delineation of tumours, segmentation of clinical images and visualization of volumes.

The next step is to segment the images to white/grey matter and extract the direction of white fibres using data with Diffusive Tensor MRI. As a subsequent step, the model uses numerical methods to predict the concentration of glioma after some time. Parameters for therapy are also included in the model, thus differentiating the results for different schemes.

The predicted tumour is then visualized in either two or three dimensions, with the initial state of the tumour being visible. An example of visualizing a simulated brain tumor after 65 days of evolution, using the DoctorEye Tool, is presented in Figure 1. The initial tumour is reddish, while the simulated tumour is greenish. The modelling and therapy parameters used in the diffusive models are initially extracted from the literature. The model is then validated using clinical data from the same patient taken at different times. Real and simulated data are compared by clinicians, and metrics of model efficiency are extracted.

Techniques

The core idea behind the implementation of our diffusive model is that concentration of glioma cells in brain follows the diffusion reaction equation. This partial differential equation cannot be solved algebraically, thus iterative numerical methods have been developed for approximating its solution. Finite Differences have been used for meshing the brain, while the biconjugate gradient method solves the derived big sparse system of equations. The result of the solver is the estimated concentration of glioma cells in the brain after a specific period of time.

In future, ContraCancrum aims to incorporate the deformation of brain matter caused by proliferating cells. Biomechanical Engineering techniques suitable for this purpose have been proposed in the literature and are currently under research.

The partners of the project include the medical branch of Philips Technologies GmbH (Germany), Saarland University (Germany), University College London (United Kingdom), University of Bedfordshire (United Kingdom), University of Bern (Switzerland), University of Karlova (Czech Republic), Foundation for Research and Technology – Hellas (Greece) and the National Technical School of Athens (Greece).

Links:
The ContraCancrum Project:
http://www.contracancrum.eu

DoctorEye Tool:
http://biomodeling.ics.forth.gr/

Please contact:
Kostas Marias
ICS-FORTH, Greece
Tel: +30 2810 391 672
E-mail: kmarias@ics.forth.gr

Reliable Pathways Toward Multiscale Modelling

by Karl N. Kirschner, Axel Arnold, and Astrid Maaß

Multiscale modelling requires the transferral of knowledge gained at different resolutions. Through the use of an expert-driven workflow we have developed reliable pathways for transferring information from quantum mechanics to atomistic and coarse-grained simulations.

The research field of computational chemistry had its beginnings in the early 1900s. However, its widespread use wasn't realized until the last half of the century, coinciding with significant increases in computer power and algorithm sophistication. A useful way to categorize the field is by considering resolution, in terms of both particle size and molecular event time scale. For example, examining the behaviour of electrons interacting with nuclei is best accomplished using quantum mechanics (QM) calculations, while understanding the binding of a ligand to a biological receptor is easily done using atomistic molecular dynamics (AT-MD) simulations. In the last two decades there has been an increased interest in studying systems at multiple resolutions, which is commonly known as multiscale modelling. One of the strengths of multiscale modelling is its ability to provide and link a system's physics at different time scales, thus providing a unique understanding that is otherwise difficult to obtain.

One of our goals is to model chemical and biological systems at different levels of resolution, ranging from QM calculations to coarse-grained MD (CG-MD) simulations as illustrated by Figure 1. A challenging task for all multiscale modelling is to transfer the knowledge gained from one resolution to another. As such, there is an ongoing need for the development of knowledge transfer procedures.

Consider, for example, our desire to model lipid layers. QM calculations give us insight into the molecular potential energy surface of small lipids. We can also divide lipids into different chemical functionalities (eg hydrocarbons, alcohols) that can be explored using QM calculations, thus providing a better understanding of how these functionalities are coupled within the parent molecule. The resulting knowledge can be transferred to AT-MD simulations through force-field parameterization.

To achieve our goals, we have developed an expert-driven workflow package that allows parameter development to be partly automated. Such a workflow allows the researcher to influence how knowledge proceeds, ensuring that insights gained from QM are effectively transferred to atomistic simulations. The following is a brief outline of the workflow:

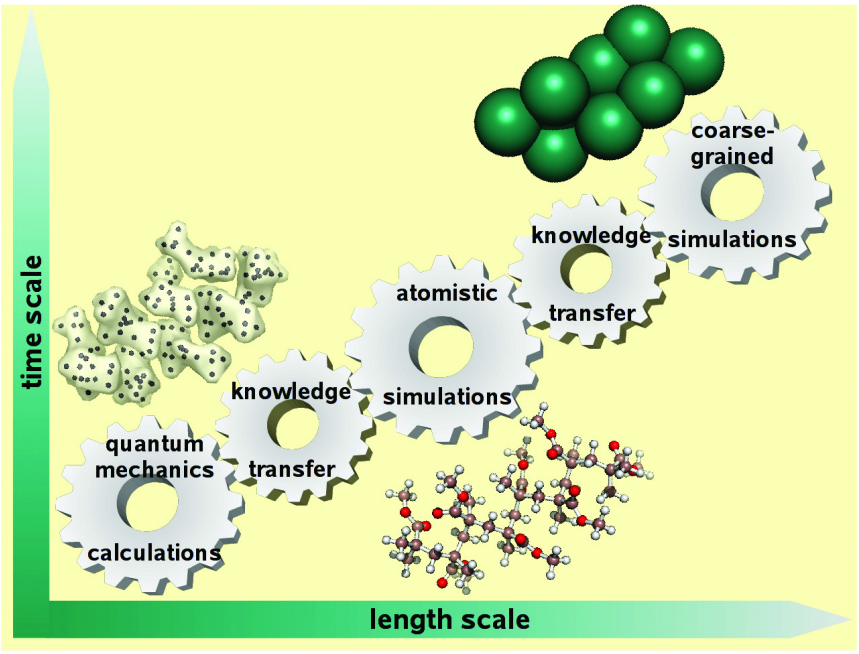

Figure 1: A representative scheme for multiscale molecular modelling. Poly(methyl methacrylate) polymer, or Plexiglas, is shown as a representative molecule.

1. Reliable QM geometry optimizations are performed, constraining the degree of freedom of interest.

2. A single self-consistent field calculation is performed at a more rigorous QM level (eg including electron correlation) on each optimized geometry, providing more reliable relative energetics.

3. Selected conformations are chosen based on the molecule's potential energy surface; subsequently, additional QM calculations are performed to determine appropriate partial atomic charges.

4. Atom types are assigned to each atom. Molecular mechanics (MM) optimizations are performed, constraining the same degree of freedom as in Step 1.

5. The parameter that corresponds to this degree of freedom is optimized through comparison of the QM and MM relative potential energy curves (ie Steps 1 & 4).

The researcher influences the workflow by deciding, in part, which small molecules are relevant for parameterization (eg alcohols for terms involving the hydroxyl group), which QM levels are appropriate, which conformations are most relevant for computing partial atomic charges, and which conformations are preferentially weighted during parameter optimization. The transfer quality of the QM-gained knowledge to AT-MD simulations ultimately resides in the expertise of the researcher performing the transfer.

Using our workflow package, we have developed a reliable force field called Explicit Torsion Parameters (ExTrM), which uses QM data for parameter optimization and experimental data validation. Our goals for ExTrM are that (i) it is transferable between molecular systems; (ii) only explicit parameter terms are used; (iii) the number of atom types is kept to a minimum, (iv) the parameters are kept as small and close to a chemical meaning as possible; (v) 1-4 nonbonded and electrostatic scaling factors are unused; and (vi) gas- and solution-phase properties can be reproduced. Currently, we have performed over 6000 QM calculations to optimize nearly 200 parameters, covering a wide range of chemical functionalities. We have also shown reliable transferability and reproduction of gas- and solution-phase properties for selected small molecules.

In general, once the parameters have been optimized and validated, they are ready for implementation in an AT-MD simulation. Once finished, the atomistic knowledge gained can be transferred to a CG model. This may be accomplished by optimizing a CG potential to reproduce AT-MD observables (eg radial distribution curves). In collaboration with the Max Planck Institute for Polymer Research, located in Mainz (Germany), we are developing the Extensible Simulation Package for Research on Soft matter (ESPResSo++), which is specifically designed for simulating CG models, as well as the necessary tools to transfer the atomistic-gained knowledge to this resolution.

Using the above techniques, we are poised to deliver molecular insight at different resolutions for many molecular and biological systems. Such systems include, but are not limited to, organic compounds, ionic liquids, solvents, drug-like compounds, carbohydrates and lipids. Possible insights include important molecular interactions, thermodynamic properties, density, diffusion rates and toxicity.

Links:
Fraunhofer Institute for Algorithms and Scientific Computing (SCAI):
http://www.scai.fraunhofer.de/en

Max Planck Institute for Polymer Research:
http://www.mpip-mainz.mpg.de

ESPResSo++:
http://www.espresso-pp.de

Please contact:
Karl N. Kirschner
Fraunhofer SCAI, Germany
Tel: +49 2241 14 2052
E-mail: kkirsch@scai.fraunhofer.de

Modelling Complex Systems with Statistical Mechanics: The Computational Approach

by Pierluigi Contucci, Cristian Giardinà, Claudio Giberti and Cecilia Vernia

Real-world phenomena are often described by complex systems with competitive and cooperative behaviour. Such systems, as much as the described phenomena, are hard to understand in a scientific perspective mainly due to the lack of general exact solutions. For cases like this, the computational sciences provide a very useful virtual laboratory. The case of disordered systems is an example of scientific computing techniques being used to test theoretical predictions and uncover new phenomena that remain unreachable by traditional analytical methods.

Complex systems made by a large number of interacting components have been studied using simplified mathematical models. A typical and paradigmatic case is the spin glass (a magnet with imitative and anti-imitative random interactions) with a rich and nontrivial structure of multiple equilibrium states. Although even the basic models are far from being solved there have been many important results, such as Parisi's analytical solution of the mean-field Sherrington-Kirkpatrick model (SK), in which every agent interacts with all the others with the same average strength. The solution has shown that the model behaviour can be conveniently described in a compact way by a function which encodes the complex hierarchical organization of the system. Such a breakthrough has influenced the hard sciences in the last thirty years no less than Onsager's solution (for the two-dimensional Ising model) did in the past. Combinatorial optimization (with the device of survey propagation algorithm) is among those fields that have been changed by its far-reaching consequences.

The broad generality of the spin-glass problem can be understood in a simple sociological context: consider a group of people where friends and enemies are present. When making a choice, friends tend to imitate each other while enemies tend to take opposite decisions. Individuals will experience conflicting requests and frustration due to the impossibility to satisfy all the relational constraints. There will typically be many sets of choices that minimize the total discomfort. That picture gives an idea of why simulations of spin glasses are difficult to perform. First, the problem is interesting for large numbers of individuals and second, the exploration of the entire set of optimal solutions belongs to the NP-complete class, making it a computationally hard problem (see Link 1 below).

A research group consisting of people from the Universities of Bologna, Modena and Reggio Emilia, Eindhoven and Roma has been involved in spin-glass computational study since 2003. The research project started with the development of a class of new optimization dynamical algorithms (see Link 2) which, using a suitable annealing procedure coupled with a balanced greedy-reluctant strategy drive the system towards the deepest minimum of the cost function. More recently the group started a computational study of the Edwards-Anderson (EA) model with short-range interaction and no known exact solution. The EA model is among the most interesting and unsettled problems in spin-glass theory, and the relevance of the mean-field solution to the EA model is vividly debated. Numerical investigations can be extremely useful to discriminate between the different physical pictures that have been proposed. The aim of the

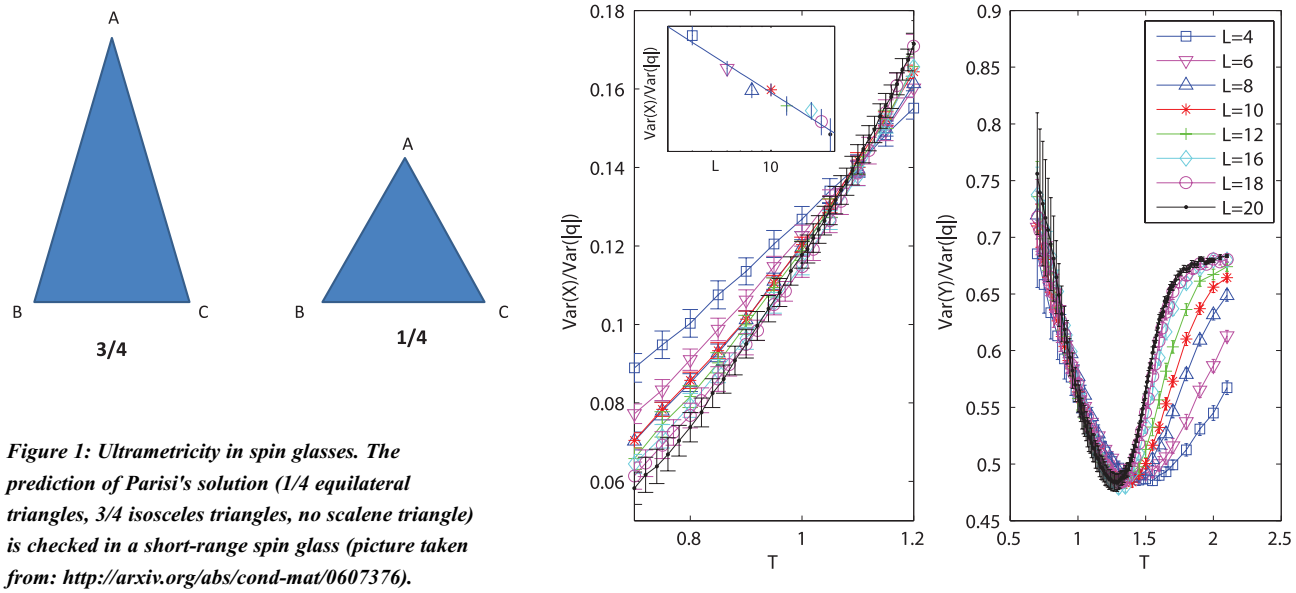

Figure 1: Ultrametricity in spin glasses. The prediction of Parisi's solution (1/4 equilateral triangles, 3/4 isosceles triangles, no scalene triangle) is checked in a short-range spin glass (picture taken from: http://arxiv.org/abs/cond-mat/0607376).

research has been to check whether and to what extent the EA model displays the features of the SK model, such as factorization properties, overlap equivalence, ultrametricity and decay of correlations (see Link 3). Ultrametric spaces have a striking property: the triangles (see Figure 1) constructed by joining three points have at least two equal sides and scalene triangles do not exist.

Scientific computing allows the behaviour of moderately large systems to be investigated by a suitable extension of conventional Monte Carlo methods. In the above-mentioned papers, the so-called Parallel Tempering algorithm was used. A well-performing implementation of the procedure, called Multi-Spin Coding, could be executed on general-purpose parallel computers. The outcome of this research gives strong support to the mean-field picture for EA spin glass. However, it is worth noting that interpreting the simulation outcomes is usually not trivial and only hints at the rigorous results. Most of all it calls for a further and deeper study of the behaviour of the realistic models of spin glasses.

Links:
1) http://www.claymath.org/millennium/ P_vs_NP/

2) http://arxiv.org/abs/math/0309058
http://arxiv.org/abs/math-ph/0309063
http://arxiv.org/abs/math-ph/0407078
http://arxiv.org/abs/0807.1197

3) http://arxiv.org/abs/cond-mat/0503155
http://arxiv.org/abs/cond-mat/0510663
http://arxiv.org/abs/cond-mat/0607376,
http://arxiv.org/abs/0902.0594

Please contact:
Pierluigi Contucci
Bologna University, Italy
E-mail: contucci@dm.unibo.it

Cristian Giardinà, Claudio Giberti,
Cecilia Vernia
University of Modena and
Reggio Emilia, Italy
E-mail:
cristian.giardina@unimore.it,
claudio.giberti@unimore.it,
cecilia.vernia@unimore.it

Engineering Chemical Substances via Molecular Simulations Utilizing Efficient Gradient-Based Optimization Algorithms

by Marco Hülsmann, Thorsten Köddermann, and Dirk Reith

The Fraunhofer Institute for Algorithms and Scientific Computing (SCAI) has developed a software tool for the automated parameterization of force fields for molecular simulations using efficient gradient-based algorithms. This tool, combined with well-established simulation techniques, can quantitatively determine many physicochemical properties for given compounds.

In the field of process engineering performed by chemical and pharmaceutical industries, physicochemical knowledge like vapor-liquid and liquid-liquid equilibria, or understanding the relationship between molecular structure and macroscopic properties, is essential. In the past two decades, the availability of increasingly powerful computers has opened routes towards obtaining various structural, dynamic and thermodynamic properties from relatively inexpensive computational models.

Molecular simulations relate macroscopic phenomena to their roots in molecular interactions. The practical applicability of these simulations in process engineering requires the construction of appropriate molecular models for a wide range of chemicals. In this context, the Fraunhofer Institute SCAI has developed the GRadient-based Optimization Workflow (GROW) software tool, which makes possible the fast and reliable development of models. For their customers, SCAI is now able to calculate quantitative physicochemical properties of chemicals and their mixtures at a wide range of temperatures and pressures. Furthermore, it is possible to explore in silico how changing the molecular structure influences chemical properties – thus guiding future experiments. With GROW, SCAI is well prepared to assist industry in its endeavour to engineer new chemical substances.

Molecular simulations are widely used to support the development process of new materials. Through simulations researchers are able to predict qualitative trends quite well. However, the key to quantitative property predictions is the accuracy of a simulation's foundation, the force field. While the equation's functional form is usually straightforward, the force field parameterization is often tedious. Manual adjustment and optimization is at best extremely time-consuming. In our pursuit to create tailor-made models for specific investigations in a timely fashion, an automated parameterization scheme is therefore essential.

While intramolecular model parameters and partial atomic charges can be obtained from quantum mechanical calculations, the determination of intermolecular parameters, eg Lennard-Jones parameters, is much more challenging. The aim is to fit these parameters to a selection of experimental properties, such as density, enthalpy of vaporization, self-diffusion coefficient, vapur pressure and reorientation time. The resulting force fields can be used to

reliably predict other properties of pure components as well as heterogeneous solutions.

GROW is a program tool kit that facilitates the gradient-based numerical optimization of force field parameters. Its components include various efficient optimization algorithms, analysis scripts and I/O handling. At the core of GROW is a gradient-based minimization of a quadratic loss function. The loss function is taken between experimental and simulated physicochemical properties at various temperatures. Specifically, the minimization algorithms implemented in GROW are the Steepest Descent, Newton, Conjugate Gradient and Trust Region methods.

Through an automated iterative process, in a parallel computer environment, the optimized solution is realized and new molecular models are developed. GROW can be coupled to various standard simulation engines, making it a powerful companion for many application fields.

GROW is flexible with regard to the simulation engine, the molecules of interest and the physicochemical properties used for target fitting. Hence, it can be broadly applied to scientifically and industrially relevant problems.

Experimental target values can be typically matched within a few percent over the considered temperature range for common properties; for example, the error is approximately 0.5-1% for density and 2-3% for enthalpy of vaporization and self-diffusion coefficients. This has been illustrated for carbon disulfide, benzene, phosgene and ionic liquids, and in general, fewer than ten optimization steps are required to obtain results in this order of magnitude. The physicochemical properties that can be simulated include shear viscosity, self-diffusion coefficients, electrical and thermal conductivity, density, compressibility, enthalpy of vaporization, vapor pressure, octanol-water partition coefficient, solubility of compounds, activity coefficient and surface tension.

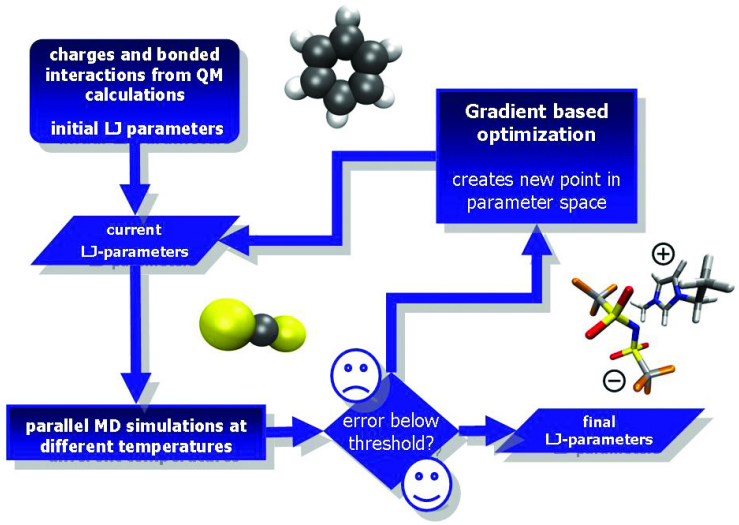

Figure 1: The optimization workflow: physicochemical properties are calculated from an initial guess of the force field parameters. If the calculated properties do not agree sufficiently well with the experimental data, a gradient-based optimization is performed. This process is iteratively performed until a stopping criterion is fulfilled and the final parameters are found.

In a further step, an optimized force field obtained by GROW for a specific compound can be transferred to chemically modified structures thereof. This enables one to study the influence of the molecular structure on certain properties.

To summarize, one of the main problems in the field of computational chemistry engineering is the optimization of intermolecular force field parameters. With the development of GROW, the Fraunhofer Institute SCAI has made an important step toward solving this problem in an efficient way.

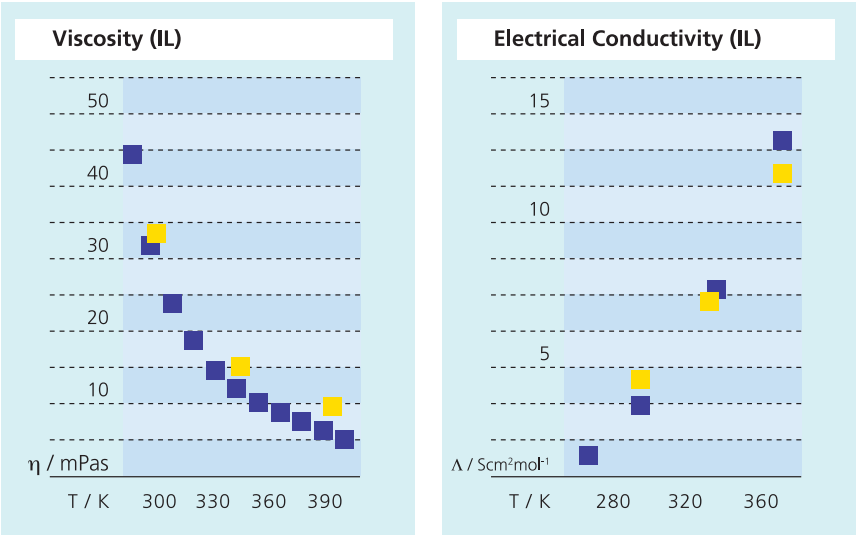

Figure 2: Simulated viscosity and electrical conductivity of an ionic liquid as a function of temperature compared to experimental data. The simulated results are in the range of the experimental error.

Please contact:
Dirk Reith
Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Germany
Tel: +49 2241 14 2746
E-mail: dirk.reith@scai.fraunhofer.de

Generic Coupling of 1D System Codes with 3D CFD Tools

by Klaus Wolf and Pascal Bayrasy

The Fraunhofer Institute SCAI has developed an application-independent interface for the coupling of different simulation codes, known as MpCCI (Mesh-based parallel Code Coupling Interface). The MpCCI interface has been accepted as a de facto standard for a neutral and vendor-independent coupling interface. Currently MpCCI supports Abaqus (© Simulia), Ansys (© Ansys Inc), Flowmaster, (© Flowmaster Ltd), Fluent and Icepak (© Ansys Inc), FineHexa and FineTurbo (© Numeca Intl), Flux3D (© Cedrat SA), MD.Nastran and MSC.Marc (© MSC Software Corp), Permas (© Intes GmbH), STAR-CD and STAR-CCM (© CD adapco), and RadTherm (© TAI). An open programming interface has been widely used to adapt customer internal codes as well as public research codes to MpCCI, thus allowing these codes to be coupled with the already supported MpCCI codes.

While typical MpCCI applications are FSI or thermal coupling, the MpCCI System Adaptor is a 1D–3D coupling link which allows the system and network code Flowmaster to co-simulate with full 3D CFD applications; combinations of 1D Flowmaster with 3D CSM (Computational Structural Mechanics) models are under preparation. This generic 1D-3D solution has been developed in order to understand how various simulation models interact with each other and what sort of impact each subsystem has on the overall system performance.

The MpCCI coupling interface combines the speed and robustness of 1D system modelling with the complexity of 3D CFD, enabling system codes like Flowmaster V7 to be used for calculations of the entire flow system and a 3D CFD to perform detailed computational calculations. For example, developing a 3D CFD model for an entire system, such as an automotive cooling system, presents significant challenges. Creating the computational models and their meshes may take a long time and the total number of cells required may make the calculations intractable (ie they may take too long, if they can be done at all). By co-simulating 1D with 3D CFD, more realistic boundary conditions and component models can be obtained, providing a deeper understanding of complex engineering systems.

Obviously this MpCCI 1D-3D combination can be used in any case where it is necessary to look at the detailed flows and designs of particular components within networks. For example:

- virtual prototyping of gas turbine blades: CFD to model the interblade cavities, CSM to model the blades and 1D to model the in-blade cooling channels
- engine cooling systems: modelling the cylinder head and cylinder cooling-jacket with CFD and the external cooling circuit with 1D
- fuel systems (1D) coupled with carburettors (3D)
- within engines (or any combustion process): fuel (1D) + air (1D)
- combustion chambers (3D) - exhaust (1D).

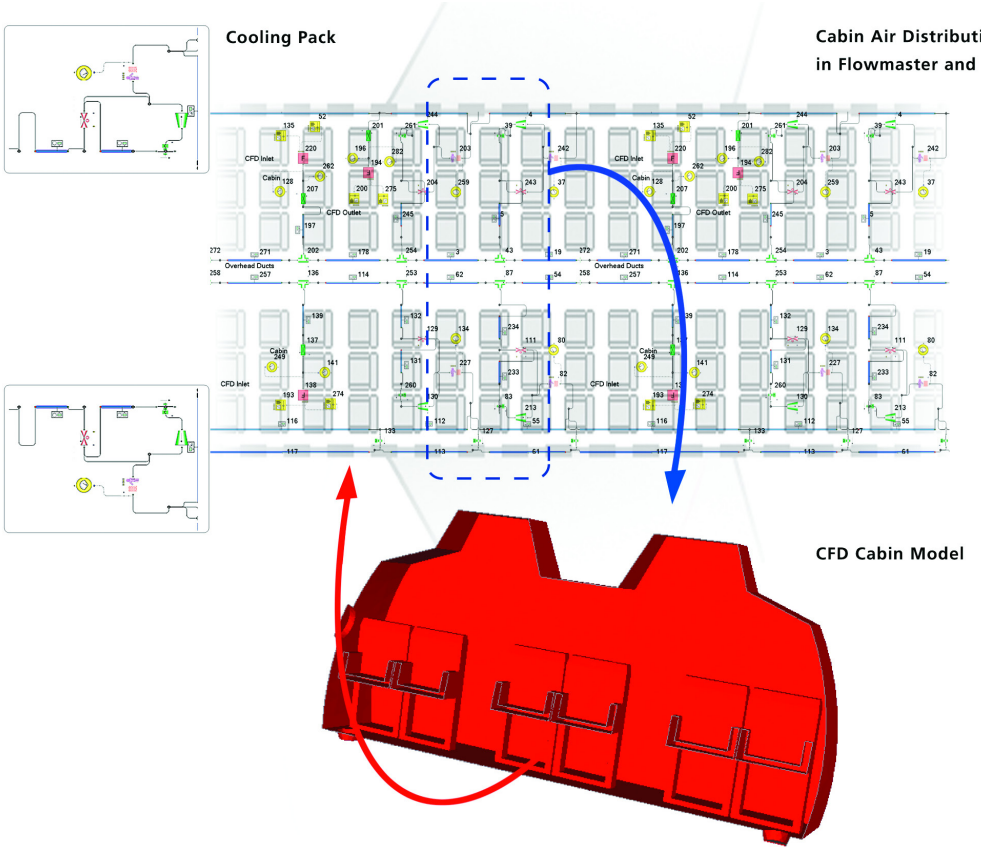

Figure 1: Coupled Simulation of the environmental control systems (ECS) and the cabine internal flows in a mid-size aircraft.

A concrete example of the use of MpCCI is provided by Environmental Control Systems (ECS) in aircraft, which are designed to optimize passenger comfort by providing satisfactory cabin pressurization, temperature and humidity control, whilst minimising risks to passenger health from airborne toxins and diseases.

In a first study it was investigated how the cooling load discharge from an ECS system affects the flow behaviour inside a typical mid-size, wide-body aircraft passenger cabin. The investigation uses 1D software Flowmaster to model the cooling pack, cabin air distribution system, and three quarters of the passenger cabin. The centre section of the cabin is modelled using 3D CFD package (Fluent or STAR-CCM), with the co-simulation middleware MpCCI providing coupling adapters to ensure that two-way, bilateral exchange of boundary parameters between the 1D and 3D CFD models gives continuity of mass and momentum transfer.

Initially the co-simulation server attempts to procure a converged solution with the 3D CFD model. Resultant boundary pressures and/or flow rates at the interface to the 1D model are exchanged with the 1D model so that it may update its solution (generally a quick process) and exchange and update the boundary conditions for the 3D CFD model. There are two 'controls' that the user may use to influence overall convergence behaviour. These are the number of iterations of the 3D solver before a data exchange of boundary conditions, and a relaxation factor which controls the extent of changes between successive calls to the 1D solver. When these parameters are appropriately specified, convergent behaviour is generally observed.

Link: http://www.mpcci.de

Please contact:
Klaus Wolf
Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Germany
E-mail: klaus.wolf@scai.fhg.de

Stochastics and Computation for Turbulent Systems

by Daan Crommelin and Jason Frank

The multiscale character of many natural systems poses a major challenge for computational studies. Often the variables of interest are macroscopic: researchers care about phenomena on large spatial scales and long timescales, not about the details of the microscopic (small-scale) behaviour. However, if the small scales do not merely perturb the large scales but fundamentally alter their behaviour, it becomes impossible to simulate macroscopic behaviour without taking into account microscale influences. Resolving the microscales explicitly usually requires such high model resolution that it becomes computationally infeasible to do sufficiently long simulations of the macroscopic behaviour (or even to simulate the macroscales at all). One approach to this problem is to use stochastic methods to represent the small scales, thereby making the macroscale simulations feasible.

The question of how to represent unresolved scales is well known in the study of turbulent systems such as the atmosphere-ocean system and other systems involving fluid flow. 'Microscopic' processes such as small-scale turbulence or cloud formation have a major impact on 'macroscopic' (planetary-scale) phenomena in the atmosphere and oceans. For explicit modelling of atmospheric convection processes and cloud formation, a model resolution of better than 50 metres is required. However, models for weather forecasting and climate studies have resolutions that are orders of magnitude coarser than this: the state-of–the-art operational model at the European Centre for Medium-Range Weather Forecasts (ECMWF) was upgraded to 16 km horizontal resolution as recently as January 2010. The atmospheric components of climate models, because they are used for simulations over much longer timespans (years to decades, or

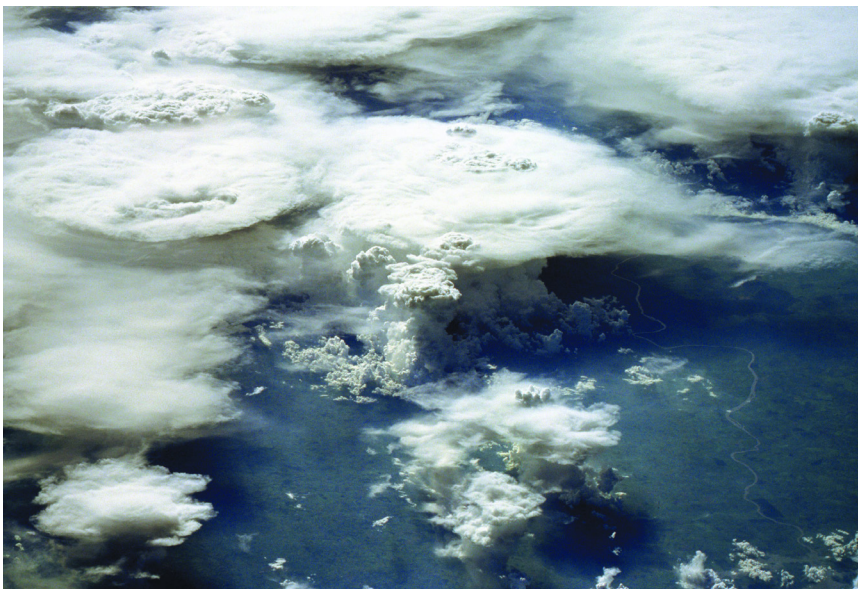

Convection in cumulus clouds is the major source of vertical transport of heat and moisture in the atmosphere. Reflection and absorption of solar radiation by clouds is important for the global heat balance. Explicit modelling of convection requires 50 metres model resolution, whereas global atmosphere models used in climate science have resolutions of 100 kilometres and more. (Image by NASA-JSC).

more), have horizontal resolutions that start at 100 km.

Another area in which this question is highly relevant is molecular dynamics. Straightforward simulations of conformation changes of large molecules (which are important to understand the chemical properties of these molecules) are slowed down enormously by the need to account for fast, small-scale motions and vibrations of the molecules. These fast motions are not particularly interesting in themselves, but they can trigger changes in the conformation of the molecules.

Formulating a good representation of (collective/aggregated) microscale influences on macroscopic, coarse-grained model variables is a major challenge. In the group Dynamical Systems and Numerical Analysis at CWI Amsterdam, stochastic approaches to meeting this challenge form an important research theme. The feedback of microscales into macroscales is often non-deterministic: a given macrostate can elicit different responses from the microscales at different times; this can be due to chaotic dynamics or inherent stochasticity at the microscopic scales. This uncertainty calls for a stochastic representation of the microscopic feedback.

Among the various activities at CWI relating to this topic, one line of research concerns so-called dynamical thermostats. Thermostatting is a well-known technique in molecular dynamics studies, but its use in fluid-dynamical simulations is rather new. The idea is to view the resolved flow model as being in contact with a reservoir of subgrid scale vorticity. Instead of modelling the unresolved flow itself, the thermostat models the energy flux between the resolved and underresolved motions in such a way that the equilibrium statistical mechanics of the resolved dynamics is preserved. In this way a statistically consistent numerical closure is obtained.

Another approach being developed at CWI aims to represent microscale influences locally in space by Markov processes which are conditioned on resolved model variables. Microscopic phenomena are represented by networks of stochastic processes (Markov chains), coupled to differential equations for macroscopic variables. The coupling goes both ways: micro-to-macro and vice versa. The properties of the Markov chains are estimated from data. This interplay between stochastics, statistics and dynamics provides a new and promising computational strategy for the study of multiscale systems. It is currently being applied to the representation ('parameterization') of atmospheric convection in a recently commenced collaborative project involving CWI, the Royal Netherlands Meteorological Institute (KNMI) and Delft University of Technology.

Links:
http://www.cwi.nl/en/research-groups/Dynamical-Systems-and-Numerical-Analysis
http://homepages.cwi.nl/~jason/
http://homepages.cwi.nl/~dtc/

Please contact:
Daan Crommelin, Jason Frank
CWI, The Netherlands
E-mail: Daan.Crommelin@cwi.nl, J.E.Frank@cwi.nl

PRO-CHAIN: Efficient Statistical Analysis of Process Chains

Tanja Clees and Daniela Steffes-lai

Among the results of the Fraunhofer project CAROD (Computer-Aided Robust Design) is a novel strategy for the statistical analysis and multi-objective robust design-parameter optimization of chains of production processes. This strategy, PRO-CHAIN, is built upon several software tools that allow for an efficient sensitivity, stability and robustness analysis, even for simulation results on highly resolved grids. Within CAROD, concrete results have been obtained for a highly crash-relevant part of a car. In this case, the strategy also includes new material and damage models and comprises both physical experiments and numerical simulations.

Important material and process parameters, geometry and also external influences can vary considerably during the fabrication of products. These variations can have a substantial and even critical influence on the robustness of production processes and the quality of the resulting products. Analysing how variations influence processes and possibly minimizing them belongs to the most challenging research and development tasks today.

This is especially true for the consideration of whole process chains. Considerably better forecasting quality of numerical simulations can be achieved by including the history of the process in the analysis of the last process step, which is still commonly considered separately. Exemplary applications in automotive engineering are the 'forming to crash' and 'casting to crash' process chains (both considering first the process of building up a part of a car, then the process of crashing the part), and in semiconductor simulation the 'process to device to circuit simulation' process chains.

A novel strategy and accompanying software tools (DesParO, DIFFCRASH) have been developed for the statistical analysis of stability, sensitivity and robustness as well as multi-objective robust design-parameter optimization of process chains. The strategy consists of the following analysis steps:

- process step 1 (eg forming)
  - stability analysis of the simulation model [optional]
  - parameter sensitivity analysis and efficient reduction of the design space: a small number of design points comprises the first 'design-of-experiment' (DoE) which can be iteratively extended if necessary in order to handle nonlinearities efficiently

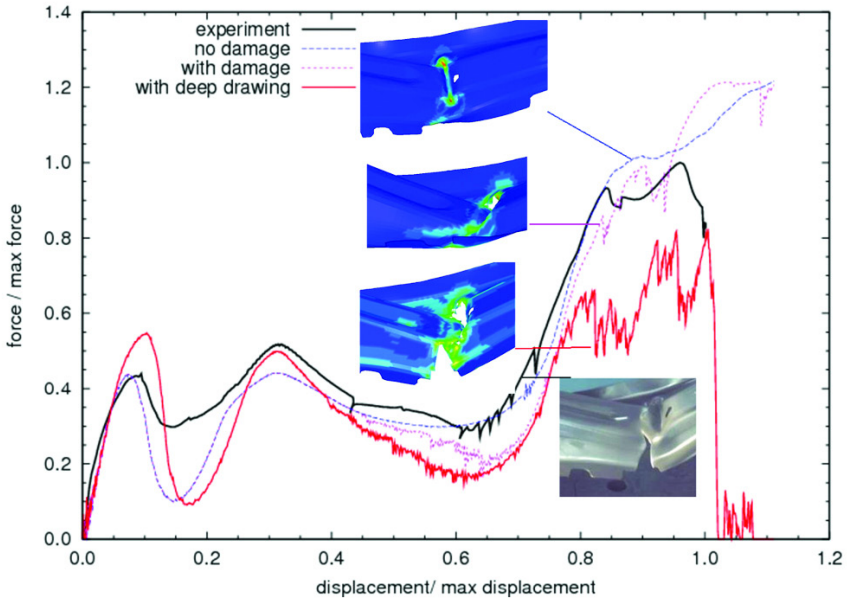

Figure 1: Physical experiment and simulations in comparison.

- robustness analysis or multi-objective robust optimization [optional]

• transfer
  - preparation and compression of the database
  - mapping (interpolation/restriction of functionals on grids)

• process step 2 (eg crash)
  - stability analysis of the simulation model [optional]
  - parameter sensitivity analysis and efficient reduction of design space
  - robustness analysis or multi-objective robust optimization of the whole chain.

The software tools can handle huge databases efficiently, even on highly resolved grids. Large random fields are therefore directly analysed, which is a prerequisite: highly resolved simulation grids are used per step of the process chain; additionally, data must be transferred (interpolated/restricted) from step to step.

During CAROD, the Fraunhofer institutes IWM and SCAI worked together on analysing the 'forming to crash' process chain for an exemplary real B-pillar, in general a highly decisive part for the crash safety in case of a side impact. More precisely, the largest out of four blanks of the B-pillar was considered. It consisted of micro-alloyed steel (ZStE340). The process of forming the sheet starting from a plain blank and a subsequent crash experiment for the formed blank were analysed in detail.

Here, the database for the analysis comprises not only forming/crash simulation results for thicknesses and plastic strains, as is commonly the case, but also damage information. Including the latter turned out to be a crucial point. The database could be reduced to 10% of its original size (on top of 'per-result' compression!) without losing important information. Comparisons of experiments with simulation results (see Figure 1) clearly show the advantages of SCAI's PRO-CHAIN together with IWM's experimental facilities and the novel so-called BI-FAILURE damage model: a considerably increased forecasting quality of numerical simulation is obtained by considering damage information from the forming step as well as variations of thicknesses, strains and damages caused by parameter variations. In particular, PRO-CHAIN allowed an intense nonlinear interplay between a crack and a kink (see Figure 2) to be detected and its relationship to parameter variations characterized.

Links:

http://www.carod.fraunhofer.de (Web site of the Fraunhofer project)
http://www.scai.fraunhofer.de/robust-design (including links to DesParO and DIFF-CRASH)

Please contact:

Tanja Clees
Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Germany
Tel: +49 2241 14 2983
E-mail: tanja.clees@scai.fraunhofer.de

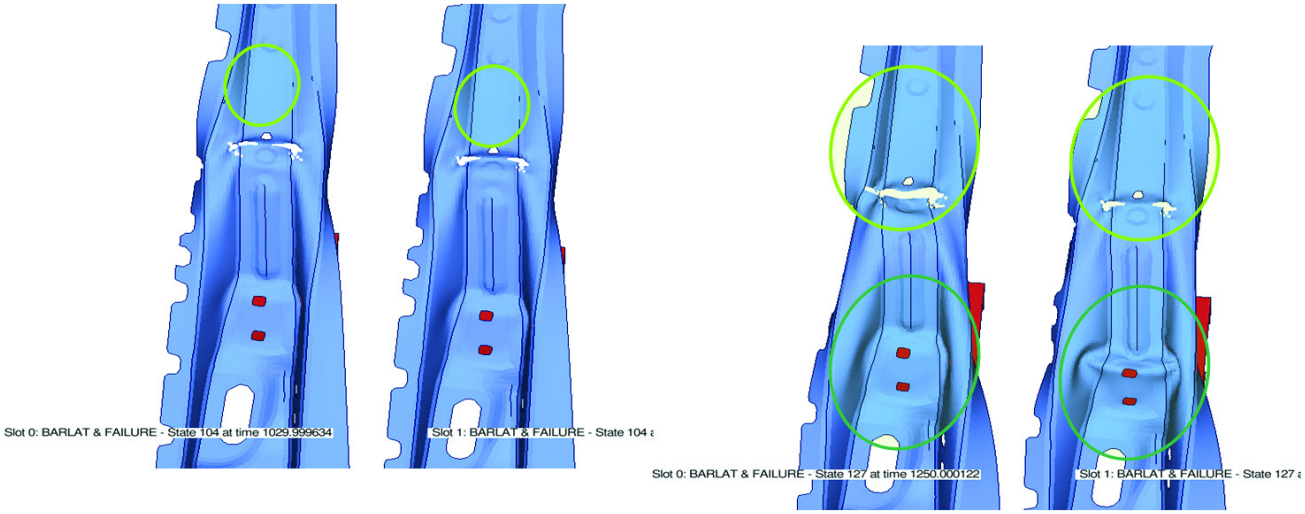

Figure 2: Interplay of crack with kink in an early (left) and later (right) time step. Shown are extreme simulation results, caused by parameter variations.

Nonlinear Time Series Analysis: A High Performance Computing Approach

by Ismael Marín Carrión, Julio José Águila Guerrero, Enrique Arias Antúnez, María del Mar Artigao Castillo and Juan José Miralles Canals

Physicists and computer scientists from the University of Castilla-La Mancha are currently performing transdisciplinary work on nonlinear time series analysis. This research will analyse some important properties of time series, and will also provide a set of high performance algorithms that will allow this analysis to be made in a reasonable time, especially in real applications such as biomedicine or climate science in which real-time responses are required.

Many applications of science and engineering, eg in physics, biology, economics or meteorology, are determined by dynamical systems. These systems evolve across time and then generate a set of data spaced in time called time series. The analysis of time series from real systems, in terms of nonlinear dynamics, is the most direct link between chaos theory and the real world. Information that is very useful for making predictions over dynamical systems can be extracted by the analysis of these time series.

Since many of these applications must provide a real-time response, it is necessary that analysis and prediction be performed on a reasonable time scale. High performance computing gives a feasible solution to this problem, which enables it to be solved in an efficient manner. Nowadays, parallel computing is one of the most appropriate ways of obtaining important computational power.

This work is included in a transdiciplinary framework, within which are collaboratively working physicists, medical doctors and computer scientists from Real-Time and Concurrent Systems (RETICS) and the Interdisciplinary Research Group in Dynamical Systems (IRGDS) from the University of Castilla-La Mancha. This framework is intended to achieve an algorithmic solution of high performance for the computational problems that underlie nonlinear time series analysis.

We have explored several algorithms of nonlinear time series analysis, emphasizing the algorithms for computing the embedding dimension and the estimation of largest Lyapunov exponent of a dynamical system. These parameters play an essential role in the identification of chaos and prediction in time series data. The embedding dimension provides us with an approximate order of the dynamical system under study. From this parameter, we obtain the maximal Lyapunov exponent that is of special interest because it provides the horizon of predictability of the system, that is, the time within which we can make good predictions.

From a computational point of view, there are some algorithms in the literature that have been developed for computing these parameters, as well as many libraries. One of the most popular and widely used for the scientific community is the TISEAN package. It includes sequential implementations of methods of nonlinear time series analysis.

However, as far as the authors know, there exist no parallel implementations of the aforementioned methods. Thus, we have developed different parallel implementations based on distributed memory architectures, taking the TISEAN package as a starting point. The parallel implementations use the message passing paradigm, and particularly the MPI (Message Passing Interface) standard. This paradigm can be extended to shared and distributed-shared systems, and therefore the developed implementations can run in many

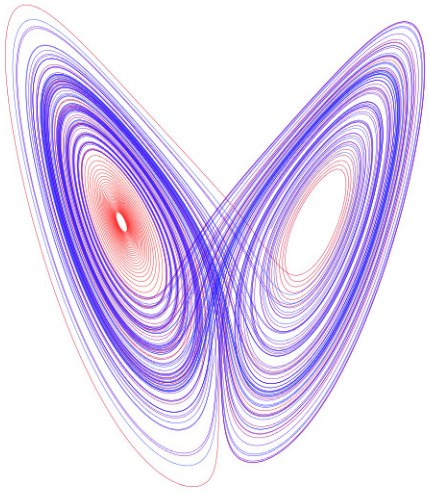

Figure 1: The Lorenz attractor.

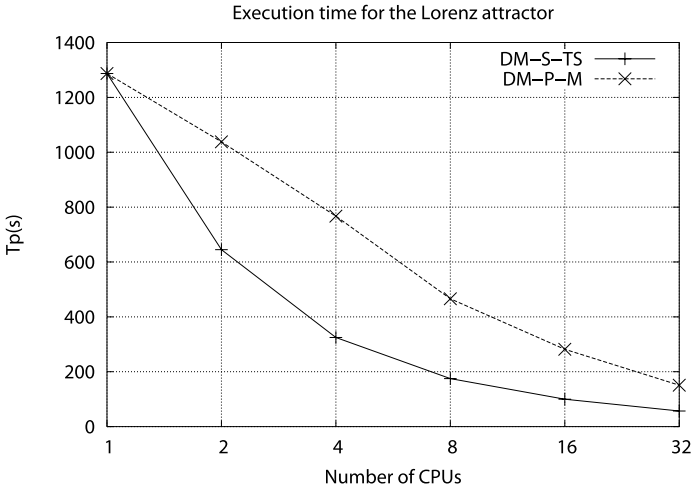

Figure 2: Execution time of the parallel implementations for computing the method of Kantz.

diverse parallel computers. MPI basically provides interfaces to send/receive data and synchronize operations between the multiple tasks of a parallel application.

Three classical examples of discrete- and continuous-time chaotic systems, the Lorenz attractor, the Hénon map and the Rössler attractor, and two real cases, an ECG signal and a temperature time series have been considered as case studies in order to evaluate the goodness of the parallel implementation. The performance analysis of the parallel implementations carried out has shown that the execution time for applying both methods has been dramatically reduced because of the use of parallelism.

The work presented here has two natural continuations. The first one consists in developing parallel implementations based on MPI+OpenMP. While often used in scientific models for shared memory parallelism on symmetric multi-processor (SMP) machines, OpenMP can also be used in conjunction with MPI to provide a second level of parallelism for improved performance on clusters having SMP compute nodes. Programs that mix OpenMP and MPI are often referred to as hybrid codes. Thus, this kind of parallel implementation is able to exploit the parallelism present in most modern parallel computers. Other future work is to drive the parallel implementations based on distributed memory architectures towards the framework of GPU (Graphics Processing Unit) architecture. The second natural continuation of this work consists in the study and analysis of complex networks and the development of appropriate software that can deal with real problems.

Please contact:
I. Marín Carrión, University of Castilla-La Mancha, Spain
Tel: +34 967 599 200
E-mail: ismael.marin@uclm.es

I-DARE: Intelligence-Driven Automatic Reformulation Engine

by Antonio Frangioni and Luis Perez Sanchez

The I-DARE system aims at helping practitioners to bridge the gap between mathematical models cast in their natural form and the myriad of available specialized solvers capable of exploiting the valuable, but possibly hidden structures in the model. It does this by automating the search for the best combination of (re)formulation, algorithm and parameters (comprising the computational architecture), until now a firm domain of human intervention.

Complex, hierarchical, multi-scale industrial and natural systems generate increasingly large mathematical models. Practitioners are usually able to formulate such models in their 'natural' form; however, solving them often requires finding an appropriate 'reformulation' to reveal 'structures' in the model which make it possible to apply efficient, specialized approaches. The search for the 'best' formulation of a given problem, the one which allows the application of the algorithm that best exploits the available computational resources, is a painstaking process requiring considerable work by highly skilled personnel. Experts in solution algorithms are required to figure out which pair (formulation, algorithm) is best, considering issues like the appropriate selection of the several obscure parameters of each algorithm.

We aim to improve the average effectiveness and efficiency of this search in the (formulation, solver, configuration) <f,s,c>-space by developing a software system capable of automating its main steps. The I-DARE system is divided into three main parts: Front-End (FE), Core System (CS) and Solving Section (SS).

The FE part comprises one or more graphical and/or textual FEs for the system. A general Modelling Environment Handler (ME-Handler) declares all functionalities that I-DARE exposes to each modelling environment, setting the interface between the FE and the CS. The interface relies on I-DARE(im), a logic-based intermediate modelling language designed using Frame Logic (Flora2), which offers high deduction power with which to analyse (query) the models.

The CS is further subdivided into three components: formulation and reformulation, performance evaluation and control. The formulation component, denoted 'Structures and ARRs', is responsible for the definition of the structured model and of the corresponding structured instance, together with the set of reformulation rules. I-DARE(lib) is the package containing all the structures that I-DARE knows. A formulation complemented with actual data becomes an 'Extended Model'; the Instance Handler (I-Handler) allows data to be retrieved from different sources. A general deduction system, I-DARE(t), can then be used to reformulate EMs by using the database 'ARR' of Atomic Reformulation Rules.

The performance evaluation component, denoted 'ML', is responsible for predicting the performances of each algorithm (considering its parameters) on a given reformulation of a given instance. This is a difficult task, for which different general-purpose and/or specialized Machine Learning approaches can be used, all implementing an abstract General Machine-Learning Control (GMLC) interface. Provided that the possible solvers for a given structure are not too many, an effective ML approach would provide the tools for performing the search in the <s, c> subspace, leaving 'only' the (re)formulation space to be explored. Other fundamental subcomponents of this component are the Continuous Learning mechanism which updates the

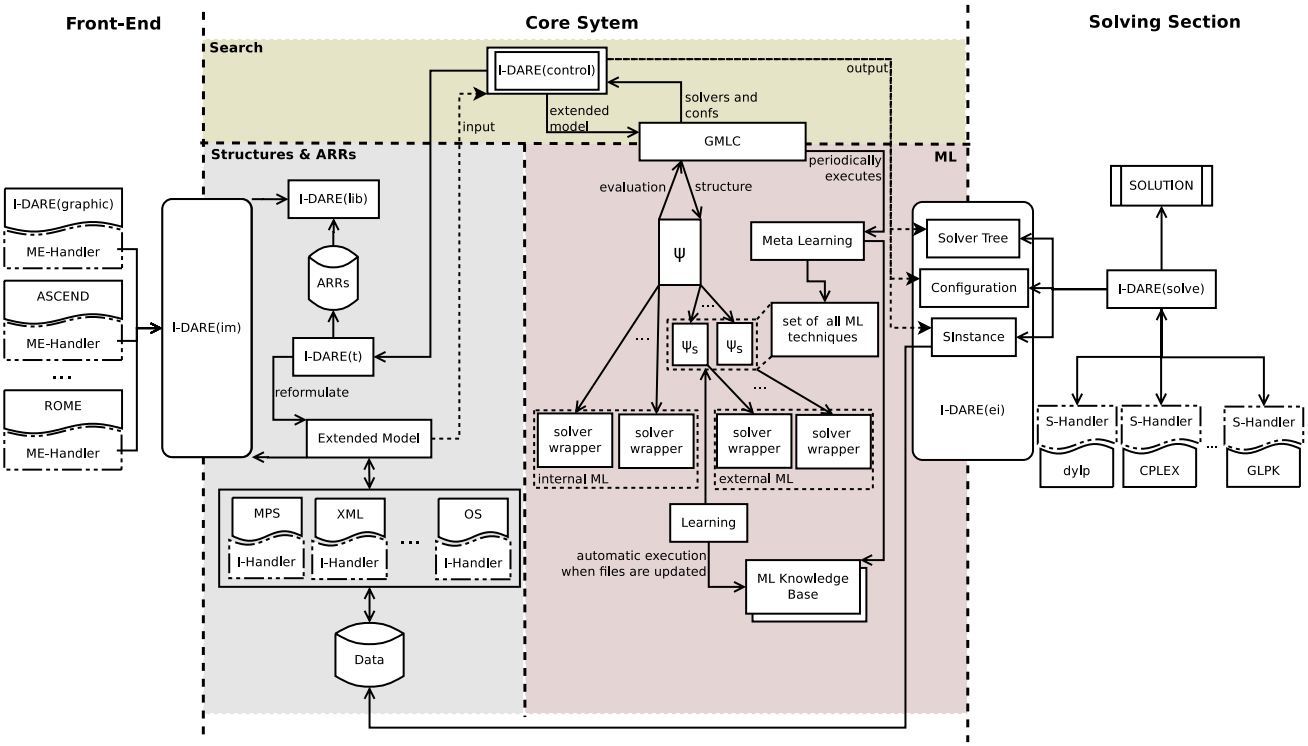

Figure 1: I-DARE architecture.

knowledge base using all the outcomes of the solutions of the system, and the Meta-Learning mechanism that makes it possible to evaluate and compare different ML techniques.

The control component is composed of the package I-DARE(control), in charge of guiding the search for the best <f,s,c> triple. This package defines the abstract interface that any control mechanism will have available to guide the deduction process of I-DARE(t) for generating the tentative reformulations; the performances will be predicted by GMLC, until the desired <f,s,c> is reached.

The final part of the I-DARE system, the Solving Section, is responsible for applying the solution method on the chosen (re)formulation, and collecting and presenting the results to the CS, which will in turn refactor them in the format of the original instance to present them to the FE. Its interface with the CS is the I-DARE Enhanced Instance format, itself composed of three <f,s,c> parts. The I-DARE(solve) package will activate the process relying on the available set of Solver Handlers (S-Handler), ie implementations of the general solver interface that allows specific solvers to be plugged into the I-DARE system.

The implementation of the I-DARE system means dealing with deep and challenging issues: the development of a general formal definition of a notion of 'best reformulation' capable of expressing the different forms that are useful to practitioners, comprising those that cannot be obtained by application of simple syntactical rewriting rules; the development of a language that can be used to formulate large-scale 'structured' models and the reformulation rules that allow one model to be transformed into a different one; the development of practical algorithms for searching the <f,s,c>-space; the design of a general interface for numerical solvers which is capable of accommodating and exploiting structural information; and finally, the implementation of all the above into a coherent and functional software system, providing proof that an industrial-strength system is eventually possible.

Tackling these challenges can potentially have profound, lasting and disruptive effects on many facets of the development and deployment of mathematical models and the corresponding algorithms to solve them. While at the beginning the system targets a narrow set of structures, mostly coming from decision and optimization problems, it is conceptually open to integration of different sets of mathematical components from almost all fields of human speculative and practical activities. Proving that it is possible to streamline systems and to automate the exploitation of scientific and technological advances through the use of sophisticated ICT tools has the potential to bring about a true scientific breakthrough.

Links:
http://compass2.di.unipi.it/TR/Files/TR-09-13.pdf.gz
http://compass2.di.unipi.it/TR/Files/TR-09-18.pdf.gz

Please contact:
Antonio Frangioni
Dipartimento di Informatica,
Università di Pisa, Italy
E-mail: frangio@di.unipi.it

Luis Perez Sanchez
Dipartimento di Informatica,
Università di Pisa,
perez@di.unipi.it

Integrating Rule-Based Modelling and Constraint Programming for Solving Industrial Packing Problems

by Abder Aggoun, Nicolas Beldiceanu, Mats Carlsson and François Fages

Packing items in bins is an old but nevertheless challenging combinatorial problem with numerous applications in industry. We report on an original approach based on constraint programming and rule-based modelling, which has been investigated in the framework of the FP6 'specific targeted research project' Net-WMS (Towards integrating virtual reality and optimization techniques in a new generation of Networked businesses in Warehouse Management Systems under constraints). It has applications in the automotive industry.

Existing Warehouse Management Systems (WMSs) provide advanced features for managing the movement of items within warehouses, but fail to comply with increasing demand for more numerical handling. Generally, WMSs lack optimization functionalities and advanced packing tools for determining how to pack items on a pallet, how many cartons are needed to pack customer items, how to pack pallets in a truck according to stability constraints and delivery route plans, and on a larger scale, how to redesign a storage area, an assembly line or similar.

These combinatorial optimization problems are particular instances of the classical 'bin packing' problem, which has contributed to the development of the theory of algorithms and complexity since the early 1970s. The one-dimensional bin packing problem (1BP) is, given N items of possibly different lengths to pack in K bins of a given length capacity, to determine whether a packing solution exists (decision problem), or to determine the minimum value K* of K for which there exists a solution (optimization problem). 1BP directly models simple forms of truck loading or assembly-line design problems. It is however an NP-hard problem, which means that no efficient polynomial-time algorithm exists for solving all instances of this problem, if we admit the conjecture P≠NP.

Two-dimensional (2BP) and three-dimensional (3BP) bin packing problems model many loading problems, respectively by layers or directly in the three dimensional space. These problems contain many variants according to the shapes allowed for the items (squares, rectangles, polytopes etc), and whether or not items may be rotated. All these variants have the same theoretical complexity (NP-hard) as long as a grid of integer coordinates and a finite number of rotations are considered for the packing. The core container loading and pallet loading problems are 3BP problems, but while in bin packing the only objective is to achieve high volumetric use, the real problems require additional constraints.

Rule-based modelling

When packing objects, it is not enough to pack them efficiently. Objects must be packed correctly. For instance, heavy objects must not be piled on top of fragile ones. Boxes must not be left hanging partly unsupported, and so on. Any business that needs to pack something has regulations about how objects may and may not be packed. Unfortunately such knowledge is typically written on paper, in natural language, and is therefore difficult for computers to access.

A key challenge of our work was to come up with a way of encoding this knowledge in a form that a computer can process and use for optimization, and that a human without a computer science degree can read, understand and edit. In this way, not only can the computer come up with packing plans that are space efficient, but also guarantee that such plans obey the business packing regulations.

Most people find it easier to understand or describe a complex design in small pieces than as a monolithic whole. Similarly, in computer science, there exists a formalism for describing knowledge in small pieces: rules. So we decided to define a rule-based formalism, or language, for business packing regulations. Moreover, it must be possible to transform knowledge expressed in our language into executable code for tasks like computing packing plans.

```
container_loading_constraints(Items, ItemsReferences, Bin, BinSize, Dims) -->
         domains(Items, BinSize, Dims) and
         lexicographic(ItemsReferences, Dims) and
         non_overlapping(Items, Dims) and
         gravity(Items) and
         stack_oversize(Items, 10) and
         stack_support_area(Items, 100) and
         stack_weight_sum(Items) and
         weight_balancing(Items, Bin, 1, 10).

gravity(Items) -->
         forall(O1 in Items,
           origin(O1, 3) = 0 or
           exists(O2 in Items, O1:uid # O2:uid and on_top(O1, O2))).

stack_weight(Items) -->
         forall(O1 in  Items,
                  forall(O2 in Items,
                           (O1:uid # O2:uid and above(O1, O2))
                           implies
                           lighter(O1, O2))).
```

Figure 1. Example of PKML rules defining container loading constraints, gravity and stacking rules.

Initially, we aimed at defining a rule-based language specifically for modelling packing problems. Later, we realized that limiting the scope to packing was no advantage, and widened it to include modelling general discrete optimization problems. The resulting language was called Rules2CP, or Rules To Constraint Programming. Rules2CP allows the definition of library modules. We defined such a module tailored to modelling packing problems and named it PKML. It provides a set of generic primitives for use in packing rules.

All optimization software developed in the project is based on the constraint programming approach. The project uses two constraint programming platforms: the open-source Choco Java and the commercial SICStus Prolog platforms. A significant effort went into the development of Rules2CP compilers: one compiler into Java for Choco, and the other into Prolog for SICStus.

We also designed and implemented an alternative approach: to compile a subset of Rules2CP into side-constraints for the 'geost' constraint, which is a computational workhorse in our optimization software. The point is to achieve higher efficiency from a tight integration of multiple logical conditions in one single constraint.

Geometrical constraint solving

Constraint programming is a declarative programming paradigm which relies on two components: a constraint component which manages posting and checks satisfiability and entailment of constraints over some fixed computational domain, and a programming component which assembles the constraints of a given problem and expresses search procedures. Because of its capability to handle a great variety of specific requirements, the main technical approach developed by the NetWMS project to solve packing problems is based on constraint programming, with the aim of making significant advances on the design and implementation of efficient geometrical constraint solvers in this context.

The goal of our research was to provide a flexible geometrical kernel that could directly handle a large class of packing problems. As illustrated in Figure 2 in the context of non-overlapping, this goal has been largely achieved.

Constraint programming was selected since it offered a flexible environment in which to develop such a constraint kernel. More precisely we chose to embed the geometrical kernel within a generic global constraint, named geost, so that it could be used with all standard available constraints of a typical constraint toolkit (Choco and SICStus in our context). The geost constraint is generic in the sense that it can handle the location in space of polymorphic multidimensional objects subject to various geometrical constraints.

A first significant effort in this research was the development of a multi-dimensional sweep-based algorithm that could handle a large class of geometrical constraints in a uniform way. A second significant effort was the development of efficient methods for handling multidimensional non-overlapping constraints. In this context, general necessary conditions were developed, especially cumulative necessary condi-

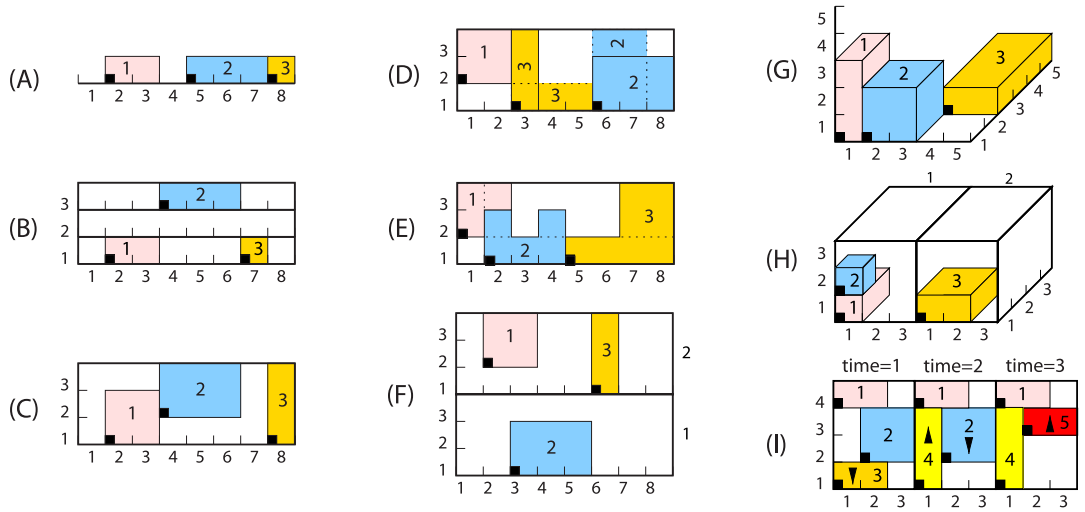

Figure 2: Typical placement problems handled by geost: Case (A) corresponds to a non-overlapping constraint among three segments. The second and third cases (B, C) correspond to a non-overlapping constraint between rectangles. (B) is a special case in which the sizes of all rectangles in the second dimension are equal to 1; this can be interpreted as a machine assignment problem. Case (D) corresponds to a non-overlapping constraint between rectangles where each rectangle can have two orientations. This is achieved by associating with each rectangle two shapes of respective sizes l x h and h x l. Since their orientation is not initially fixed, the included constraint requires that the three rectangles be included within the bounding box defined by the origin's coordinates 1,1 and sizes 8,3. Case (E) corresponds to a non-overlapping constraint between more complex objects where each object is described by a given set of rectangles. Case (F) describes a placement problem in which one must first assign each rectangle to a strip, such that no rectangle overlaps with another assigned to the same strip. Case (G) corresponds to a non-overlapping constraint between orthotopes (the generalization of a rectangle for higher dimension). Case (H) can be interpreted as a non-overlapping constraint between orthotopes that are assigned to the same container. The first dimension corresponds to the identifier of the container, while the next three dimensions are associated with the position of an orthotope inside a container. Case (I) describes a rectangle placement problem over three consecutive time-slots: rectangles assigned to the same time-slot should not overlap in time.

tions for non-overlapping. Necessary conditions for important special cases (eg pallet loading) were also conceived. Finally, in order to handle the fact that many practical problems consider only a few types of items to place, symmetry-breaking techniques that directly consider non-overlapping were designed.

The specifications of the geometrical kernel developed in this work package were implemented within both the open-source Choco Java library and the SICStus Prolog platform. A third implementation was performed outside the Net-WMS project within the open-source JaCoP constraint library.

Conclusion

By working on real-size industrial problems that take all business requirements into account, the Net-WMS project has advanced the state of the art in bin packing, constraint programming and modelling languages, making significant progress in the technological transfer toward the end users.

There are still some hard instances however, especially in 3D pallet-loading problems with rotations. In these cases the automatic placement finds sub-optimal solutions – not as good as expert solutions – and we are not aware of alternative automatic methods by which to solve these problems optimally. Current work concerns search strategies, and improvements in the geometrical constraint as well as its generalization to handle hybrid discrete continuous complex shapes.

The technological transfer of Net-WMS components is already ensured by their industrialization by our SME partner KLS OPTIM, as well as by their integration in open-source libraries or in the commercial software of our academic partners EMN, INRIA and SICS. The first version of the KLS Optimization Suite has been released, not only making Net-WMS technology available to the market at the end of the project, but also positioning Net-WMS at the beginning of a new generation of agile supply chain management systems.

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

Please contact:
François Fages
INRIA Paris-Rocquencourt, France
E-mail: Francois.Fages@inria.fr

Abder Aggoun
KLS OPTIM, France
E-mail: abder.aggoun@klsoptim.com

Nicolas Beldiceanu
EMN, France
E-mail: Nicolas.Beldiceanu@emn.fr

Mats Carlsson
SICS, Sweden
E-mail: matsc@sics.se

Modelling Gene Regulatory Networks - An Integrative Approach

by Alina Sîrbu, Heather J. Ruskin and Martin Crane

Integration of large amounts of experimental data and previous knowledge is recognized as the next step in enhancing biological pathway discovery. Here, data integration for quantitative regulatory network modelling is under investigation, using evolutionary computation and high-performance computing.

With the completion of the genome project and advances in high-throughput measuring techniques, a base for systems biology research has been created. This involves uncovering biological pathways and networks between cell products, and is an important step in finding disease markers and treatments, and in the future, toward building synthetic organisms.

Such aspirations have triggered considerable research, spanning multiple fields such as mathematics, computer science and biology, resulting in enormous amounts of data describing cellular processes, and stimulating alternative modelling efforts. However, most of these data are diverse and largely unreconciled, so that modelling approaches can offer only limited insight. Thus, in recent years, integrative approaches to modelling biological networks have appeared. However, the number of data types is usually small compared to the potential set, so that biological realism is only approximate.

The Centre for Scientific Computing and Complex Systems Modelling (SciSym) was formally established at Dublin City University in 2007. It links existing research groups in computing, such as ModSci (Modelling and Scientific Computing) and mathematics. It conducts multi-disciplinary research in complex systems modelling, ranging from biological to socioeconomic systems. One of the centre's recent projects has been to build qualitative models for gene regulatory networks (GRNs). This was funded by the Irish Research Council for Science, Engineering and Technology.
GRNs consist of interactions between proteins, known as transcription factors, and genes, which in turn encode other proteins. Transcription factors bind to a DNA region close to the target gene and activate or repress its expression, ie the formation of the encoded protein. This creates complex networks of activation and repression links which, by controlling protein levels in cells, are involved in important processes such as cell differentiation, cell cycle or response to external shock.

Two types of models have been applied to GRNs: coarse-grained, which allow for large-scale low-resolution analysis, and fine-grained, for low-scale high-resolution studies. The latter enable continuous simulation of gene expression, and are, from this point of view, important tools in predicting outcomes of different perturbations in the network, corresponding to diseases or treatments. However top-down models have the disadvantage of size limitation, with only small networks feasible for

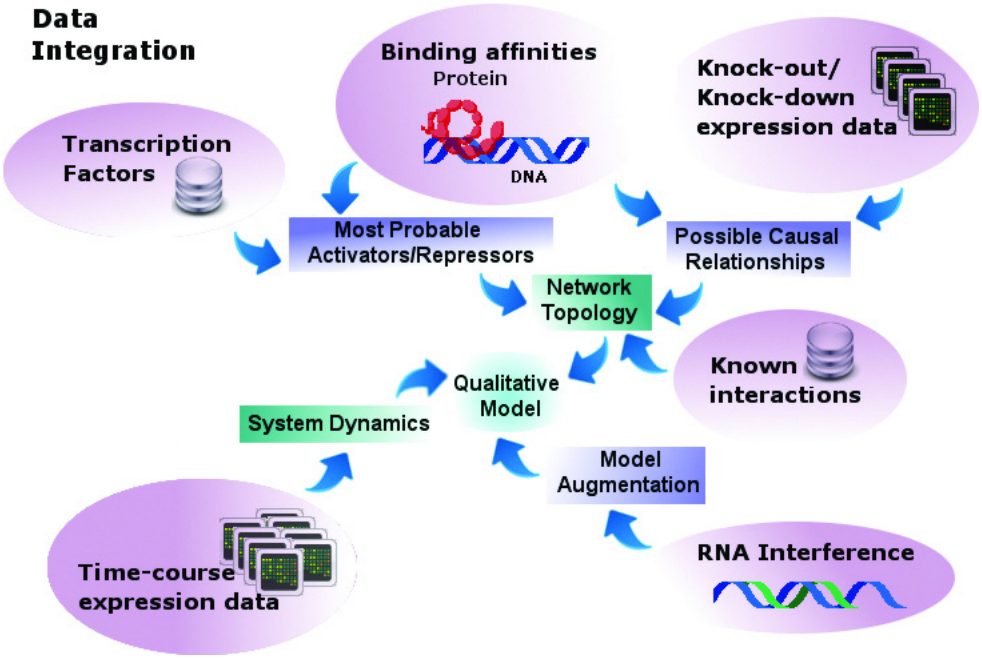

Figure 1: Data integration.

this approach and the computational power currently available. In this context, a principal aim is to increase the scale capability and quality of quantitative models, in order to perform reliable simulations of entire GRNs. From a software engineering point of view, the objective is to obtain a user-friendly software application, which handles the computational and mathematical aspects of modelling, to enable biologists to focus on validation and interpretation of results.

To achieve this objective, one solution is to adopt an integrative approach. Given the large amount of available data on gene expression dynamics in different biological processes, we investigate means to integrate these to increase the scale capabilities of quantitative gene expression modelling, using evolutionary computation. Evolutionary algorithms are known to perform very well in large search spaces and with limited data, making them well suited to this problem. They have the advantage of flexibility in terms of adding different types of data to the inferential process. Furthermore, they are intrinsically parallel, so multi-threaded implementations are straightforward, facilitating the use of the local (Sci-Sym) high-performance computing cluster, and subsequent extensions to the Irish Centre for High-End Computing facilities.

Typically, dynamical models of GRNs are inferred using time-course gene expression data. However, due to experimental costs, time series are usually very short, containing less than 100 measurements, while networks can be very large, involving hundreds of genes. There do exist multiple time series from different sources describing the same process, but they are measured on heterogeneous platforms, so analysis is not straightforward. To our knowledge, integration of these time series in the context of GRN qualitative modelling has not previously been attempted, and requires considerable pre-processing. We have analysed integrated gene expression data coming from three microarray platforms in preliminary studies, using different preprocessing methods to provide a comparative framework and the basis for a single model. Statistical integrity is a major consideration. Nevertheless, we have shown that differential equation models built from multiple datasets, are more robust to parameter and data perturbations, and display less noise overfitting. This provides the first step towards improving qualitative models for regulation.

The integration process will continue with knock-out and knock-down measurements of gene expression, introduced at the initialization stage of the evolutionary algorithm. Information on known transcription factors and interactions will be used to implement customized genetic operators, driving the inferential algorithm towards richer areas in the search space. Binding-site information will also be included in model fitness evaluation, to reward models that connect genes and proteins with high affinities. Finally, we plan to introduce RNA interference in the modelling process, first as metadata, then as distinct nodes in the regulatory network.

Links:
http://sci-sym.computing.dcu.ie/
http://www.computing.dcu.ie/~msc/
http://wiki.c2b2.columbia.edu/dream/index.php/The_DREAM_Project

Please contact:
Alina Sîrbu
Dublin City University, Ireland
Tel: +353 1 700 6747
E-mail: asirbu@computing.dcu.ie

Computational Computer Architecture Research at NTNU

by Magnus Jahre and Lasse Natvig

The computer architecture group at the Norwegian University of Science and Technology (NTNU) in Trondheim, Norway, is working on issues that are arising as increasing numbers of processors are integrated on a single chip. Discrete event simulators and high-performance computers are indispensable tools in this quest. By combining the cutting-edge multi-core simulator M5 from the University of Michigan with the 5632-core Stallo cluster at the University of Tromsø, researchers are making progress on the issues facing future multi-core architectures.

The chip multiprocessor (CMP) or multi-core architectures is a recent technological innovation that has received considerable attention in both academia and industry. The main reason is that such architectures reduce the impact of physical and economic design constraints. Consequently, a number of commercial vendors now produce CMPs, and most new desktop computers are equipped with multi-core processors. The recent popularity of CMPs is due to the following factors:

- technology scaling has made it feasible to place multiple cores on one chip
- it has become increasingly difficult to improve performance with techniques that exploit Instruction Level Parallelism (ILP) beyond what is common today
- single-core, high-performance processors consume a great deal of power, meaning expensive packaging and noisy cooling solutions are needed. This limitation is known as the power wall. For large compute clusters, reduced power consumption gives doubled benefits – reduced power demands from both the processors and the cooling systems. Consequently, multi-core computing contributes to what is now called Green IT
- when designing a CMP, a processor core is designed once and reused as many times as there are cores on the chip. Furthermore, these cores can be simpler than their single-core counterparts. Consequently, CMPs facilitate design reuse and reduce time to market.

Processor performance has been improving at a faster rate than the main memory access time for more than twenty years. CMPs do not automatically reduce the impact of this problem. In fact, they can make it worse because multiple processors need to be fed from the same slow memory. In recent years, our group has pursued several avenues towards reducing the impact of this problem. Firstly, we have proposed techniques that increase shared cache utilization. We have also looked at prefetching, which analyses the memory access stream and then attempts to retrieve data before the processor requests it. Here, we have both proposed new prefetching heuristics and illustrated how prefetching can be used to improve memory bandwidth utilization. Finally, we have proposed techniques that reduce the performance impact of destructive interference between concurrently scheduled processes running on different cores.

The Stallo Cluster at the University of Tromsø (Photo: Thilo Bubek).

Modern cycle-accurate simulators such as the M5 are complex pieces of software. Typically, they contain tens of thousands of lines of code. Most computer architecture simulators are event driven. Consequently, execution is modelled as discrete events that occur at specific times, and time is typically measured in clock cycles. The simulator is built around the time-ordered event queue. The latency of an operation is modelled by first calculating the latency and then adding an event to the event queue at the time the operation will complete. By expanding on this simple concept, it is possible to model very complex behaviour.

A computer architecture research paper can demand as much as 25 computer-years' worth of simulation. A significant part of computer architecture research is Design Space Exploration. We test our architectural techniques on a variety of different programs (called benchmarks) to ensure that the technique is sufficiently general. Then, we investigate the impact of changing key architectural features such as the amount of cache space or off-chip bandwidth. Each combination of architectural parameters and benchmarks is one point in the design space. These points are independent of all other points and can therefore be executed in parallel. Consequently, our research is well suited to large clusters. Fortunately, we have been granted access to the 5632-core Stallo cluster at the University of Tromsø by the Norwegian Metacenter for Computational Science (NOTUR). Having access to a large, professionally maintained compute cluster is a great advantage.

Links:

http://research.idi.ntnu.no/multicore
http://www.m5sim.org
http://docs.notur.no/uit
http://www.notur.no/

Please contact:

Lasse Natvig, NTNU, Norway
E-mail: Lasse.Natvig@idi.ntnu.no

Smart Agents and Sentiment in the Heterogeneous Agent Model

by Lukáš Vácha, Jozef Baruník and Miloslav Vošvrda

'Smart traders' is a new concept for improving the usual models for future price movements in financial markets. The influence of the proposed smart traders concept is examined with simulations.

An important feature of heterogeneous agent models (HAM) is their ability to explain stylized facts observed in financial time series, mainly fat tails and volatility clustering. Typically, in the heterogeneous agent's model, two types of agents are distinguished: fundamentalists and chartists. Fundamentalists base their expectations about future asset prices and their trading strategies on market fundamentals and economic factors, such as dividends, earnings, macroeconomic growth, and unemployment rates. Chartists or technical analysts try to extrapolate observed price patterns, such as trends, and exploit these patterns in their investment decisions.

We introduce a new concept – smart traders. The idea of smart traders is based on the endeavour of market agents to estimate future price movements. By adding smart traders we try to improve the original heterogeneous agent's model so it can better approximate real markets. Smart traders are designed to forecast the future trend parameter of price deviations using information sets consisting of past deviations. For simplicity, they are modelled to assume that the price deviations, defined by the model, are an AR(1) process and they use the maximum likelihood estimation method for forecasting. Thus, in our model we use two groups of traders: smart traders and a group of stochastically generated trading strategies.

Furthermore, we introduce changes in sentiment, which we define as a shift in beliefs about the future trend of a new investor strategy on the market. This allows us to model trend-followers and contrarians. In this work we use only the form of jumps in sentiment. Our main expectation is that the introduction of smart traders and changes in sentiment will change the simulated market prices significantly.

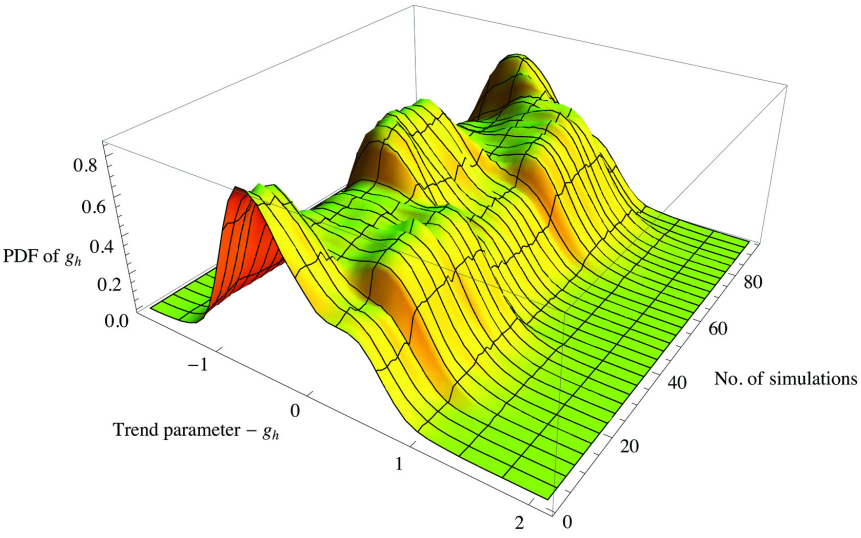

Figure 1: Empirical PDF of the trend parameter - g_h on the market through the iterations with the sentiment change.

Model

The model presents a form of evolutionary dynamics called the Adaptive Belief System in a simple present discounted-value pricing model. Simulated capital market is a system of interacting agents who immediately process new information. Agents adapt their predictions by choosing from a limited number of beliefs. Each belief (trading

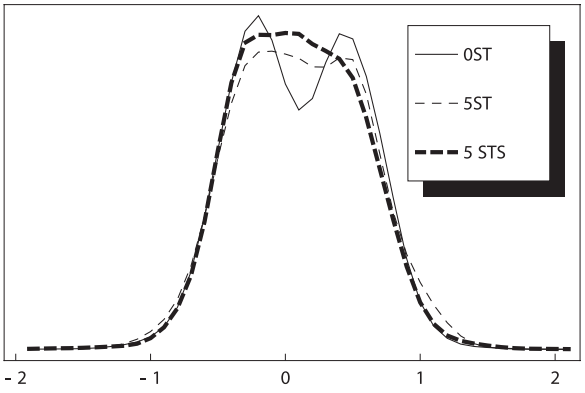

Figure 2: Empirical PDF of X(t) for simulated models without smart traders, with five smart traders, and with five smart traders and changing sentiment.

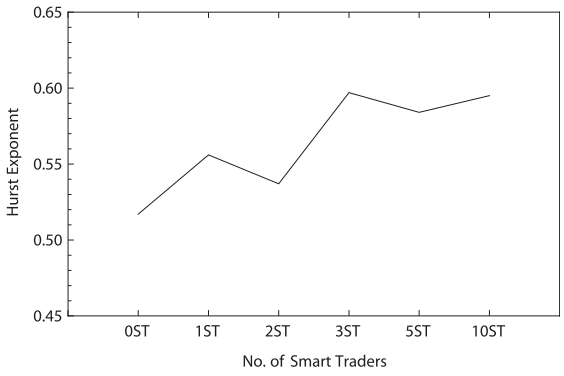

Figure 3: Hurst exponent of returns for different numbers of smart traders.

strategy) is evaluated by a performance measure. Agents on the capital market use this performance measure to make a rational choice that depends on the heterogeneity in agent information.

Simulation Results

The main purpose of the simulations is to examine the influence of the proposed smart traders concept and changes in sentiment on the simulated market prices. We compare the model without smart traders (0ST) with the model with five smart traders (5ST) and the model with five smart traders in the first group and changes in sentiment in the second group (5STS).

Altogether we consider 40 trading strategies for each simulation. Figure 1 shows the empirical probability density function of the trend parameter observed on the simulated market with sentiment change. It is the cross-section through the iterations, and the changes in sentiment can be clearly observed. Each model has been simulated 36 times to achieve robust results. Figure 2 shows the kernel estimation of the probability density functions (PDFs) of simulated returns.

Conclusion

By adding smart traders we try to improve the original heterogeneous agent's model so that it provides a closer description of real markets. The main result of the simulations is that the probability distribution functions of the price deviations change significantly with an increasing number of smart traders in the model, and they also change significantly when changes in sentiment are introduced. We also use the Hurst exponent to measure the persistence of the price deviations and we find that the Hurst exponent increases significantly with the number of smart traders in the simulations (see Figure 3). This means that the introduction of the smart trader's concept into the model results in significantly higher persistence of the simulated price deviations. On the other hand, the introduction of changing sentiment in the proposed form does not change the persistence of the simulated prices significantly.

Please contact:
Lukáš Vácha - CRCIM (UTIA)
Tel: +420-266052432
E-mail: vachal@utia.cas.cz

Holons for Analysis, Modelling and Simulation of Complex Systems

by Massimo Cossentino, Vincent Hilaire and Abder Koukam

Nobel Laureate Herbert Simon states: "Empirically a large proportion of the complex systems we observe in nature exhibit hierarchic structure." Starting from this assertion we developed a novel approach for designing and implementing complex systems combining the multi-agent approach and the holonic social organization perspective.

This project is a collaboration between the Palermo department of the ICAR institute of the Italian National Research Council (CNR), and the SeT laboratory of the University of Belfort-Montbelliard (UTBM), France. It started in 2007 and it is still running today. This collaboration was the opportunity for several exchanges, and one researcher from ICAR/CNR spent a year and a half in Belfort within this activity.

Our starting point was that in the past, researchers have deployed a considerable effort in the analysis, modelling and implementation or simulation of complex systems. Their work improved the state of the art in understanding and designing software capable of simulating such systems. However, many of these systems fail to fully represent the real system they model. The aim of this project was to provide a natural, versatile and scalable representation dimension for such systems. Such a goal was pursued by adopting the holonic paradigm in the context of multi-agent systems (MASs).

MASs naturally emphasize the importance of system properties such as autonomy, proactivity, social abilities and self-organization. These properties mean MASs are an obvious candidate for a paradigm for the design of complex systems. Indeed, this paradigm proposes new strategies for the analysis, modelling and implementation of such systems. Its elementary constituents are called 'agents', ie software entities which exhibit autonomous and flexible behaviours.

Complex systems are characterized by a large number of entities in interaction, exhibiting emergent behaviours. Nobel Laureate Herbert Simon states: "Empirically a large proportion of the complex systems we observe in nature exhibit hierarchic structure. On theoretical grounds we could expect complex systems to be hierarchies in a world in which complexity had to evolve from simplicity." Simon (1996)

This raises the question: if nature has selected this path, should scientists trying to model complex systems privilege it too? In this project we are interested in these hierarchical structures to analyse and model complex systems. More precisely, we explore the requirements and consequences of modelling complex systems using hierarchically composed agents.

In 1967 Arthur Koestler coined the term holon as an attempt to conciliate holistic and reductionist visions of the world. A holon represents a part-whole construct that can be seen both as a component of a higher-level system and as a whole composed of other holons as substructures.

We apply holonic-related concepts to systems which exhibit a kind of hierarchical structure, or which there are many types. Indeed, simulations of

complex systems often need several levels of granularity with different time/geographical scales. We applied this approach to several problems with very good results. For instance, we simulated a big industrial plant. This plant has over 19 000 employees working in different shifts to ensure the plant is running 24 hours a day. Last year over 1600 trucks entered the plant every day. Geographically, three cities and a highway enclose the plant. Such a configuration makes it impossible to simply increase the plant's size to accommodate new buildings and forces the infrastructure to be redesigned. Due to the great number of constraints and interrelated dependencies between traffic and production, a simulation tool could prove to be of great help when evaluating different designs. Even the smallest modification in a plant of this size often requires a significant budget to be invested. A reliable simulator offers the possibility of detecting 'side-effects' prior to the project's validation.

Our main contribution is the definition of a methodology, namely ASPECS (Agent Specification Process for Complex Systems; see link below), and of a development and deployment platform, JANUS. The key concepts underlying the methodology and platform are based on organizational theories. We adopted this approach because it is largely recognized in the agent research community that an MAS should be conceived in terms of an organized society of individuals in which each agent plays specific roles and interacts with other agents. As pointed out by Ferber, the organizational approach offers a number of advantages and can contribute to agent-oriented software development through heterogeneity of languages, modularity, multiple possible architectures and security of applications.

ASPECS is structured in three phases or domains. The first, the Problem Domain, provides the organizational description of the problem independently of a specific solution. The concepts introduced in this domain are mainly used during the analysis phase and at the beginning of the design phase. The second domain, the Agency Domain, introduces agent-related concepts and provides a description of the holonic, multiagent solution resulting from a refinement of the Problem Domain elements. The third and last domain, the Solution Domain, is related to the implementation of the solution on a specific platform. This domain is thus dependent on a particular implementation and deployment platform. In our case, this part of the process is based on the JANUS platform that we specifically designed to ease the implementation of holonic and organizational models.

Links:
http://www.aspecs.org/
http://www.janus-project.org/

M. Cossentino, N. Gaud, V. Hilaire, S. Galland, A. Koukam. ASPECS: an Agent-oriented Software Process for Engineering Complex Systems. International Journal of Autonomous Agents and Multi-Agent Systems (IJAAMAS). 20(2). 2010.): http://www.pa.icar.cnr.it/cossentino/paper/Cossentino_ASPECS.pdf

Please contact:
Massimo Cossentino
ICAR-CNR
Tel: +39.23842257
E-mail: cossentino@pa.icar.cnr.it

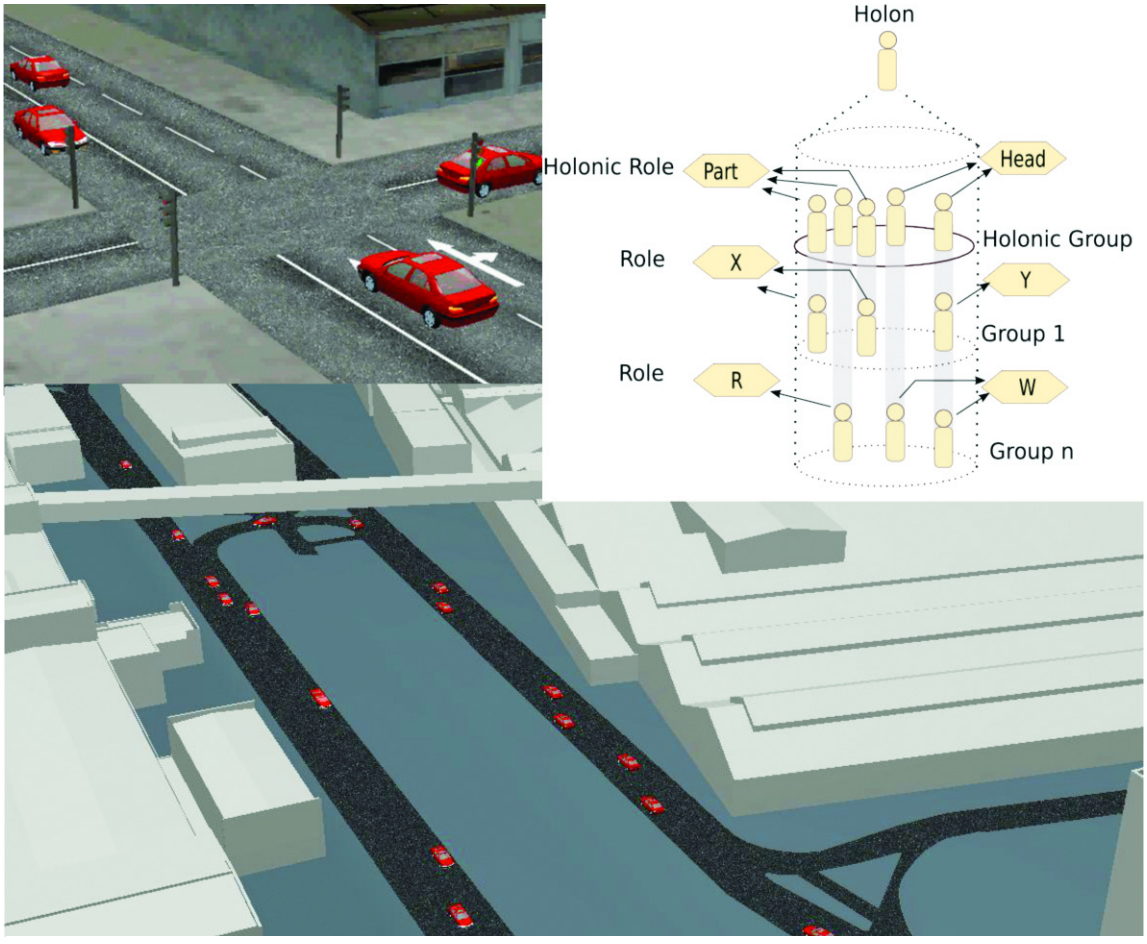

ASPECS has been succesfully employed for the simulation of a big industrial plant.

An Approach to the Parallelisation of Agent-Based Applications

by Chris Greenough, Shawn Chin, David Worth, Simon Coakley, Mike Holcombe and Mariam Kiran

The Flexible Large-scale Agent Modelling Environment (FLAME) has been developed in a collaboration between the Computer Science Department at the University of Sheffield and the Software Engineering Group at the STFC Rutherford Appleton Laboratory. FLAME is an applications program generator for agent-based simulations. Using the modeller's definition of the agent-based model and the associated C-code that implements the agent actions and state changes of the agents, FLAME generates the user's application using program templates as either a serial or parallel code.

There are many agent-based modelling (ABM) systems. Many of these systems are based on Java as their implementation language. Although a good language for Web-based and some communications applications, it is not one often used in the area of high-performance computing. Similarly there are relatively few agent systems that address the problem of scalable parallel simulations. The essence of an agent-based application is that the basic elements of the model can be represented by autonomous entities (agents) that can communicate with other agents in the model and act (perform tasks or computations or even communicate information to the agent population) based on that information. This makes them ideal to model many situations in which discrete elements interact.

Cell biology, transport logistics, behavioural science, and financial and economic markets are a few of the many application areas in which ABMs have been successfully used.

A major problem has been that these models are computationally very intensive and in a population of p agents, p^2 interactions can be common. Consequently to date only relatively small populations have been modelled.

FLAME has been developed in a collaboration between the Computer Science Department of the University of Sheffield and the Software Engineering Group at STFC Rutherford Appleton Laboratory. One of its aims was to develop implementation techniques for agent-based simulation that could exploit high-performance computing systems.

FLAME is an applications program generator for agent-based simulations. Using the modeller's definition of the agent-based model and the associated C-code that implements the actions and state changes of the agents, FLAME generates the user's application using program templates as either a serial or parallel code. This is illustrated in Figure 1: the model and agent functions are provided by the model developer and FLAME provides the program templates that the xparser uses to generate the simulation code. The final application executable is built as a serial or parallel application as required by the developer using the FLAME Message Board Library.

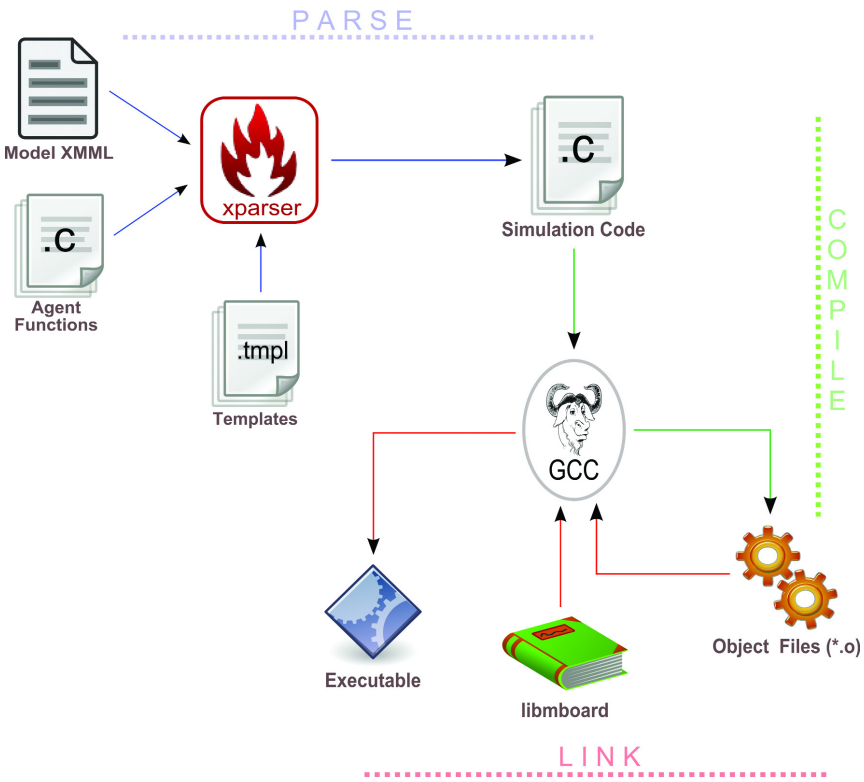

Figure 1: The FLAME framework.

It must be noted that the FLAME framework is not the application, but rather the application generator. Any approach to parallelisation must be embedded in FLAME's program templates and be applicable to a generic agent-based simulation. Although this underlying generic architecture determines much of the potential parallelism in applications, the agent model must contain that parallelism at a task level so that FLAME might be able to utilise it.

It is well understood that a knowledge of task locality and task communication are key elements in achieving efficient and scalable parallel performance. Even if an application has many potentially independent tasks that could be executed in parallel, the balance between the computational and communication loads will greatly affect any

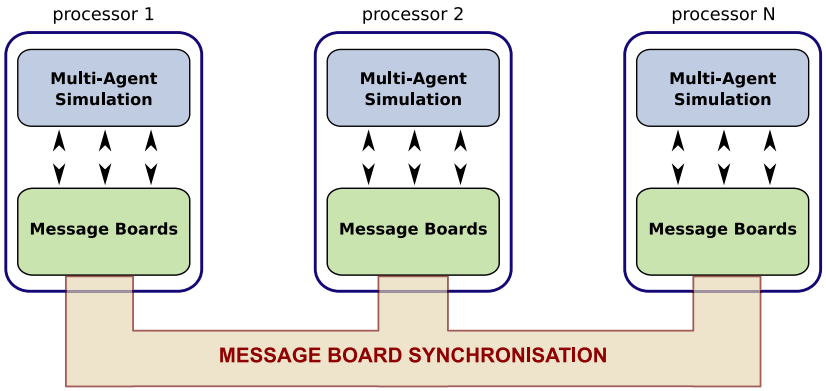

Figure 2: A representation of the parallel decomposition used in FLAME.

parallel performance. These ideas can be expressed in terms of the granularity of application: what are the relative sizes of the computational and communications loads of each of these independent tasks.

As with most computational problems there is a computational load, but significantly for agent applications there is a communications load – the information flowing between agents.

FLAME distributes the computational load – the agents themselves – over the processing elements to produce some form of computational load balance in a conventional way. However, the communications load or data load is also distributed over the processors. Each communication type between agents has its own message repository (message board) and these are distributed over the processors. During the simulation a process of message board synchronization is performed to ensure that agents receive the information they require.

FLAME combines the use of multi-threading with Message Passing Interface (MPI) to improve parallel performance. One thread is used for computational work and a second handles all communications. The placement of these message board sychronisations has a fundamental effect on the parallel performance of the application.

Each agent can be decomposed into linked state changes. These states might involve computation or communication. FLAME uses a state graph to schedule agent operations and to exploit parallelism. FLAME schedules nodes (tasks) in this graph to maximise the communication time available to any message board synchronisation task. In this way the information wait time for any agent is minimised.

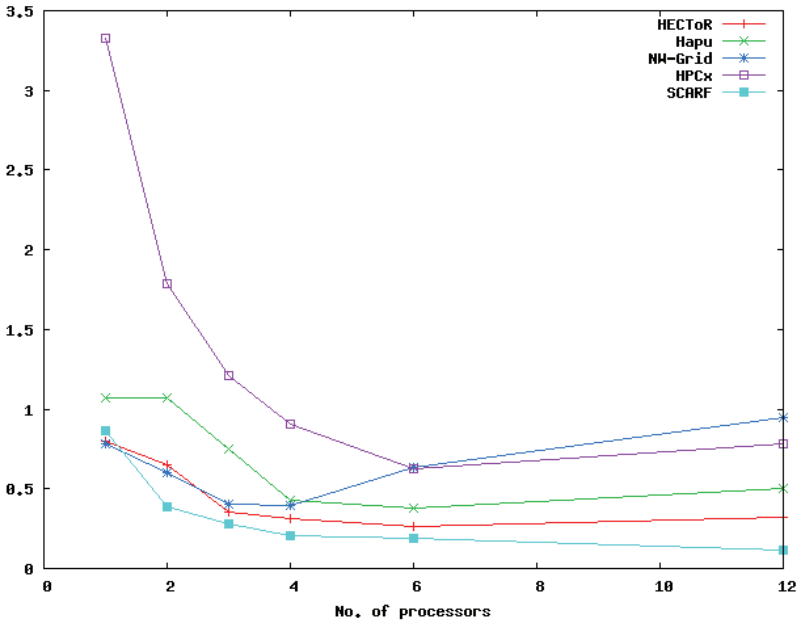

Figure 3: Some results from the EURACE model using a population of 22,000 agents.

The FLAME technology has been used with great success in a variety of application areas in both its serial and parallel modes.

Within the EURACE project, FLAME has been used to develop a very complex multi-agent simulation of the European economy involving agents such as banks, firms, households etc. With the EURACE model there were nine agent types and over 60 types of communication message. The EURACE model represented the three major markets in an economy: labour, goods and financial.

In parallel processing terms the model can be view as a tightly coupled, fine-grained application and hence not particularly amenable to large-scale parallelism. In some simpler problems from biological applications the FLAME applications have shown reasonable speedups with populations of one million agents. However, the performance of the significantly more complex EURACE model gives encouragement for the FLAME approach achieving some acceptable speedups on small numbers of processors with a population in the order of 22,000 agents.

We are undertaking further research and development to improve the communications latency in the FLAME synchronisation process, along with the development of more sophisticated message filtering techniques to reduce the data flow between processors.

This work was funded by the EURACE project (STREP-35086), which in turn is funded by the European Commission's 6th Framework Programme. Details of the partners and outputs can be found at the EURACE Web site.

Links:
EURACE: http://www.eurace.org/
FLAME: http://www.flame.ac.uk

Please contact:
Chris Greenough
STCF Rutherford Appleton Lab, UK
E-mail:
christopher.greenough@stfc.ac.uk

A Computational Approach to Patient Flow Logistics in Hospitals

by Anke Hutzschenreuter, Peter Bosman and Han La Poutré

With the aging of the population and the demand for cost-efficiency, logistics and planning in hospitals are becoming increasingly important. In many countries, hospitals are organized in a decentralized fashion, with (medical) departments and units having a high degree of autonomy in management and planning. The Agi-Care project (Agent-based intelligent health care planning) develops computational approaches to the optimization of patient flow logistics in hospitals, ie, concerning the various pathways of inpatients moving through various units in a hospital. The Agi-Care project has been carried out at Eindhoven University of Technology, in cooperation with the Catharina Ziekenhuis Eindhoven and Centrum Wiskunde & Informatica (CWI - the Dutch national research centre for mathematics and computer science), in the Netherlands.

Scheduling decisions in hospitals are often taken in a decentralized way. This means that different specialized hospital units make autonomous decisions relating to such issues as patient admissions and schedules of shared resources. Decision support in such a setting requires different methods and techniques to those described in the majority of existing literature, which tends to assume a centralized model. The design and analysis of such methods and techniques is the focus of the project, which is now in its final stages. Specifically, we developed computational models to provide dynamic decision support for hospital resource management and the prediction of future resource occupancy, and we studied the application thereof.

Hospital resource management targets the efficient deployment of resources like operating rooms and beds in various care units (intensive care, medium care, nursing room etc). Allocating resources to hospital units is a major managerial issue as the relationships between resources, their utilization, and the combined pathways of different patient groups through the hospital is complex (see Figure 1).These issues are further complicated by the fact that patient arrivals are dynamic, and treatment processes are stochastic.

Our approach to providing decision support combines techniques from multi-agent systems and computational intelligence (CI). This combination of techniques allows us to properly consider the dynamics of the problem while reflecting the existing distributed decision-making practice in hospitals. Multi-agent techniques are used to realistically model multiple hospital care units and their decision policies, multiple patient groups with stochastic treatment processes, and uncertain resource availability due to overlapping patient treatment processes (see Figure 1). Optimization and learning techniques from CI allow for designing and evaluating improved (adaptive) decision policies for the agent-based model, which can then be implemented easily in hospital practice.

In order to gain insight into the functioning of this complex and dynamic problem setting, we developed an agent-based model for the hospital care units (partly) used by the cardiothoracic surgery (CTS) department, with their different (sub)types of patients. To assess the applicability of this agent-based model, we developed an extensive simulation. Several experiments demonstrated the functionality of the simulation and showed that it is an accurate representation of the real world. The simulation is used to study decision support in resource management and patient admission control.

To further improve the quality of decision support, we studied the prediction of future hospital resource usage. Using these predictions, the future impact of taking a certain decision at a given moment can be taken into account. In the problem setting at hand, for instance, predicting the resource utilization resulting from an admission decision is important to prevent future bottlenecks that may cause the blocking of patient flow and increase patient waiting times. The methods we investigate for the task of prediction are forward simulation and supervised learning using neural networks. In an extensive analysis, we used stochastic techniques to study how accurate and precise prediction outcomes can be obtained.

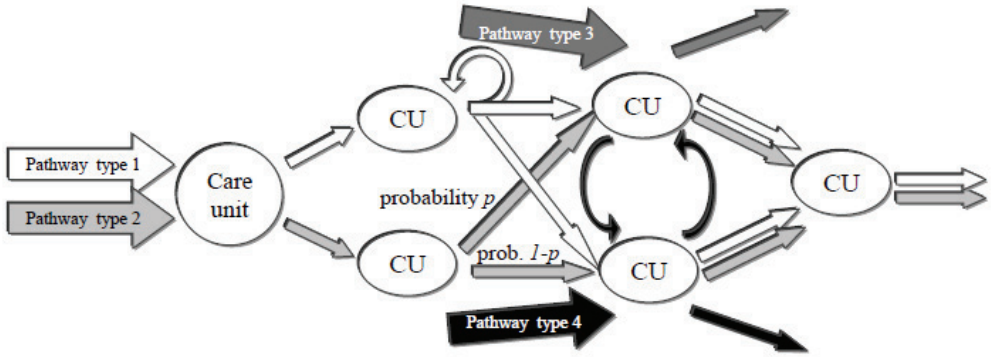

Figure 1: Example of various types of patients with their possible pathways through different care units (CU) in a hospital Every pathway type has its own color, where branchings in a pathway of a certain patient type occur with some probability, reflecting the uncertainty in patient treatment within the patient types, and where lengths of stay at a certain care unit are stochastic within a patient type.

To optimize resource allocation decisions, we considered multiple criteria that are important in the hospital problem setting. Specifically, we focused on three conflicting objectives to be optimized: maximal patient throughput, minimal resource costs, and minimal usage of back-up capacity (of beds). All criteria can be taken into account by finding decision policies that have the best trade-off between the criteria ('Pareto optimality'). We derived various decision policies that partly allow for adaptive resource allocations. The form of the policies allows them to be easily understood by hospital personnel. Moreover, we incorporated a 'bed exchange mechanism' that enables a realistic implementation of these adaptive policies in practice.

In our optimization approach, the parameters of the different decision policies were determined using a multi-objective evolutionary algorithm (MOEA). The MOEA uses the agent-based simulation in order to evaluate the quality of potential solutions, and it thus optimizes the output of the simulation (ie the three optimization criteria) as a function of the policy parameters. Optimization using the MOEA with its embedded simulations was performed on a high-performance cluster. Our results on resource management showed that the benchmark allocations obtained from a case study are considerably improved by the optimized decision policies. Furthermore, our results showed that the use of adaptive policies can lead to better results and that further improvements may be obtained by integrating prediction into a decision policy.

Anke Hutzschenreuter will defend her Ph.D. thesis that has resulted from this project in the upcoming months.

Links:
http://is.tm.tue.nl/
http://www.cwi.nl/en/research-groups/Computational-Intelligence-and-Multi-agent-Games

Please contact:
Anke Hutzschenreuter
E-mail: a.hutzschenreuter@gmail.com

Han La Poutré, CWI, The Netherlands
E-mail: han.la.poutre@cwi.nl

Computational Modelling and Simulation for an Interconnected World

by Jim Duggan

Our modern world is complex, and comprises a multitude of natural, engineered and social systems. The goal of this project is to design computational simulation approaches that can help us understand the way in which these systems interact, so that we can design better futures.

This research project, entitled Computational Methods for Modelling Complex Dynamic Systems, is a four-year research programme located in the College of Engineering and Informatics, at the National University of Ireland, Galway. Funded by Science Foundation Ireland , seven researchers are pursuing an interdisciplinary work programme to explore novel approaches to modelling complex social, economic and industrial systems. The core motivation for this research programme is to develop mathematical techniques that can help decision makers to: (i) understand the interconnected nature of our modern world; (ii) evaluate possible policy design alternatives in silico; and (iii) generate qualitative and quantitative insight into the possible future behaviours of these systems.

The project draws on three related methodological approaches for generating insight into economic and social systems:

- System Dynamics is a modelling approach based on the idea that systems can be understood by analysing their embedded feedback loops. These interlocking loop structures can be formulated as equations, and simulated using numerical algorithms. Feedback loops then can be formally analysed to demonstrate the impact of loop structures (for example, virtuous and vicious cycles prevalent in economics) on overall system behaviour.
- Agent-Based Computational Modelling is a complimentary approach to feedback modelling, where the unit of representation is an individual (or agent), and individuals have states (physical and psychological) and a range of behaviours. Individual decisions are based on an agent's state space, and information from their immediate network is also considered. The key idea of the agent approach is to generate macroscopic explanations from local behaviours. For example, in the property market, an agent-based model could represent individual consumer behaviours, and from that, generate an explanation of the recent boom and bust cycle.
- Multi-Agent Systems and Game Theory put the focus on how to combine the efforts of many autonomous agents, where each has different information and separate goals. Game theory is an essential technique for this analysis, as it provides the basis for the mathematical investigation of interactions between independent, self-interested agents. Within game theory, games can be classified as being non-cooperative or cooperative, and the application areas include economics, political science and computer science.

In the research project, these three valuable perspectives are integrated through a multi-method simulation workbench. At its core is a numerical integration algorithm, which acts on a very large set of equations, where each equation represents a state of behaviour of an individual in the population. Different agent types can be designed, and any number of these agents (within the resource limit of the machine) can be instantiated. For each simulation, an appropriate network structure is selected. The options include the classic network types such as small world, lattice, grid, random and scale free, and the user also has the opportunity to create a hybrid network structure based on any combi-

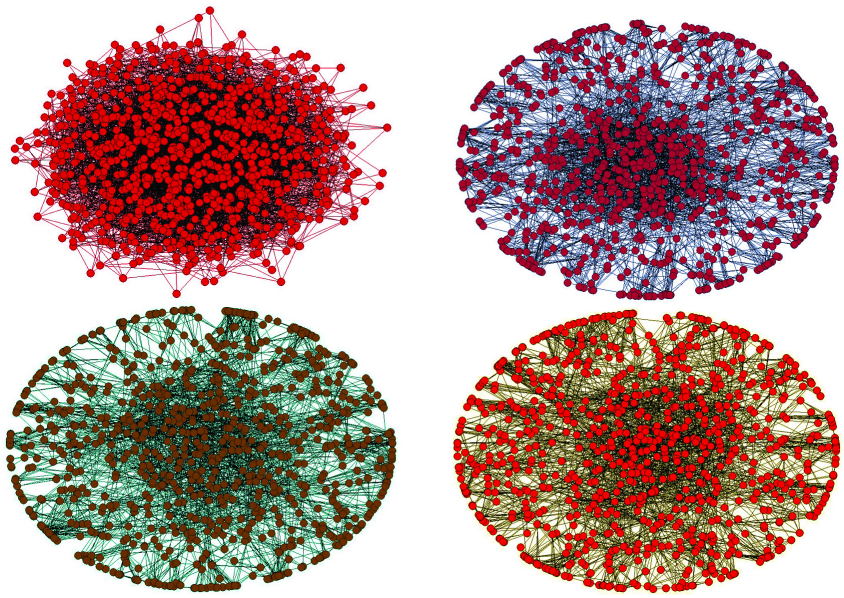

Figure 1: From random to small world – the evolution of a dynamic social network.

nation of these standard representations. Once a model has been simulated, the results can be analysed at macro and micro level, including information on agent state changes, and also how the social networks may have evolved.

Sample output from a simulation is shown in Figure 1. This visualizes the spatial structure of a network, and how it dynamically evolves over time. Starting in the northwest quadrant with a random network, in our novel network algorithm individual nodes seek out new acquaintances, and from these local actions, the overall network evolves to take on the properties of a small world network. This approach is valuable for exploring the impact of dynamic social networks on the convergence of social norms.

In order to inform the design of this workbench, the project draws on the literature of thoroughly explored and well documented complex problems. These include: the beer game, developed at MIT, which explores the impact of misperceptions in decision making across connected industrial systems; epidemiology-inspired dynamics, such as norm evolution in society, the tipping points for new product introductions, and the spread of infectious diseases; and commons-type problems where the challenge is to optimally share limited resources amongst a population of self-interested agents. The project has also identified the potential for applying these underlying computational techniques to games development, and has developed strong collaborative linkages with scientists and creative writers with a view to working in the genre known as serious games for young children and teenagers.

The project is now just past its midpoint, and a library of algorithms and computational methods has been designed and developed. To date, these have been tested on well-defined problems. Future work will concentrate on real-world case studies, and will focus on gathering empirical data (eg social networks, industrial supply chains, epidemiology data) and constructing models to help generate insight into the behaviour of complex social systems.

Link:
http://corrib.it.nuigalway.ie

Please contact:
Jim Duggan
National University of Ireland, Galway.
Tel: +353 91 493 336
E-mail: jim.duggan@nuigalway.ie

Training for CBRN-Emergencies: The Successful First Phase of the SimRad-Projects

by Gerhard Chroust, Karin Rainer and Markus Roth

In responding to the growing need to be prepared for chemical, biological, radiological and nuclear (CBRN) emergencies, First Responders (i.e. fire brigades, emergency medical services and police) must quickly evaluate such incidents and take appropriate actions to minimize negative effects on humans and goods. Achieving such interventions poses a critical challenge, since humans do not possess any inborn, natural sensors with which to recognize these dangers early enough. Additionally they are not equipped with natural, semi-autonomous reaction patterns. Nevertheless, it is of the utmost importance to avoid endangering the First Responders. This requires special training and adequate tools, especially since a considerable proportion of First Responders are volunteers providing part-time services.

The SimRad (Simulation- and information system to manage Rescue units at disaster) projects, supported as part of the Austrian KIRAS programme for safeguarding Austrian critical infrastructures, have as their prime objective the conceptualization, specification and design of models, methods and tools for training First Responders concerned with CBRN emergencies. The results of these projects can also be applied in real assignments. The first phase, SIMRAD.NBC, which finished successfully in October 2009, provides the basis for a user- and practice-oriented training

environment using state-of-the-art technologies and taking into account an up-to-date inventory of the needs of First Responders. The major issues identified and discussed were:

Characteristics of CBRN incidents and emergencies: CBRN incidents differ considerably from other incidents, due to the different nature of the origin and the behaviour of the sources of dangers (usually contaminants without directly observable menace).

Requirements for interventions in CBRN emergency situations: understanding the needs and requirements of First Responders is a key to implementing adequate support and training systems. A survey of First Responders identified and clarified prerequisites for a usable system.

Process-oriented view of interventions: an intervention is a network of processes, which may be serialized, interdependent or parallel. A process view provides detailed identification, analysis and evaluation of key processes for focused training possibilities.

Process description language: the heterogeneity and diversity of First Responders on a basic operational level requires a detailed description of the individual processes with an appropriate prescriptive language. Such a language allows systems to be built efficiently and with fewer misunderstandings, showing new ways of collaboration. It provides system builders with the opportunity to codify the knowledge of domain experts and to review the resulting systems,

Assessment of intervention processes: a consequence and benefit of the process view is the ability to assess quality aspects of individual intervention subprocesses for comparison and improvement by proposing best practices.

Organizational needs: successful interventions depend strongly on organization. Communication between different levels of hierarchy and cooperating units can become a crucial issue; gaps in communication and their circumvention were therefore investigated. According to expert opinion there is a big potential as well as a pressing need for structuring, clarifying and standardization in communication flow on these multiple levels.

SimRad framework: an effective, robust and flexible software framework is needed for training and for support of real action. It must be modular, so that various training and/or planning modules can be 'plugged in' depending on available facilities and actual and changing needs. Real-time evaluation and effortless change of these modules is necessary. In order to provide humans with flexibility, the SimRad framework is able to learn from current situations and from its users. This framework is designed to be compatible with existing devices and systems that are in use by First Responders and even by broader target groups like providers of critical infrastructure.

Protected First Responder testing contamination (Photo: SIMRAD).

Simulation of interventions (mixed and augmented realities): simulation is a proven method for training and planning (cf flight simulators) offering repeatability, controllability and the possibility for evaluation. With respect to the often invisible and for human senses unnoticeable dangers of CBRN incidents, simulation allows dangerous situations and materials to be replaced by harmless ones. Mixed and augmented Reality software plays a key role in such simulations.

Human aspects: well-trained and experienced emergency personnel are crucial to a successful intervention. A series of interviews among experts provided information on positive or negative influences (motivators and stressors) affecting the ability to take proper action. Needs and views were analysed by considering organizational and human factors, equipment and methods for deployment. For volunteers in particular, the provision of interesting and challenging tasks in training settings is a strong source of motivation.

Systemic considerations: CBRN emergency situations involve highly complex and somewhat opaque processes having wide ramifications, showing often surprising immanent dynamic behaviour. The interplay of various activities and actors needs synergistic and/or interdependent views. Different types of progression and impact of catastrophes were investigated.

SimRad.COMP, the current project, began in October 2009 and will continue until 2011. It is based on the results of SIMRAD.NBC and will elaborate the specification for an appropriate software platform and support environment. This will include the concepts for marketable products (eg simulation and/or evaluation software) that are compatible with and supported by the SimRad framework. SimRad.COMP is supported by the Austrian Federal Ministry for Transport, Innovation and Technology (BMVIT) within the KIRAS Security Research Programme.

Links:
SimRad: http://www.simrad.at
KIRAS: http://www.kiras.at

Please contact:
Peter Ziehesberger
Ziehesberger Elektronik, Austria
E-mail: peter@ziehesberger.at

Nadine Sturm
Research Institute of the Red Cross, Austria
E-mail: nadine.sturm@w.roteskreuz.at

Accurate Reconstruction of Single Individual Haplotypes for Personalized Medicine

by Filippo Geraci and Marco Pellegrini

Advanced personalized medicine is one of the goals of current research, and in this area new genetic diagnostic methods are critical. Diagnostic technology that can be used in the field (closer to the patient, and far from the traditional high-tech labs) is needed. We are studying this problem with respect to technology for personalized chromosome haplotyping.

The Human Genome project (1990-2003) highlighted the fact that humans share over 99% of their DNA. In the last few years however, the attention of researchers has shifted from what is common among members of the same species (the genome), to what is different. The single nucleotide polymorphisms (SNPs, pronounced 'snips') are the most common form of variation in human DNA. The total number of SNPs present in the human genome is estimated to be about 9-10 million and currently about 3.1 million have been profiled. The set of SNPs present in a chromosome (called the haplotype) is of interest in a wide area of applications in molecular biology and biomedicine. Personalized haplotyping of (portions of/all) the chromosomes of individuals is one of the most promising basic ingredients leading to effective personalized medicine (including diagnosis and therapy).

Luckily, personalized haplotyping is now becoming feasible due to a steady decrease in the cost of sequencing equipment that will soon render cost-effective the effort of producing complete genomic profiles of individuals versus collecting data on predefined genetic markers. The first publication of a complete individual diploid human genome sequence appeared in 2007. The current aim of the research community is to be able to sequence individual genomes at a cost of just 1000 USD within a few years. The actual trend in the cost of state-of-the art technology suggests that this target is realistic.

A key component of the reconstruction pipeline is the so-called 'Single Individual Haplotyping' problem (SIH). PCR-based shotgun technology produces a large collection of overlapping fragments, each containing a few SNPs. By using reference maps it is relatively easy to locate the position and orientation of such fragments in the chromosome, but not to determine the association of the fragments with the two homologous copies (paternal and maternal) of a chromosome. From the biological point of view, the single individual haplotyping problem consists in determining the association of each fragment with one of the two homologous chromosomes and then determining the haplotype.

From the computational point of view the problem was first formalized by Giuseppe Lancia in 2001. It can be described as follows: each fragment is represented as a string of fixed length such that each character corresponds to a nucleotide or to a gap (if the fragment does not cover a certain position). Each SNP can be covered by a certain number of fragments and can take only two values (the values of the haplotype in that position). The natural way of representing fragments is to store them in a matrix (the SNP matrix). We say that two fragments (rows of the SNP matrix) are in collision if they have a different value for at least one SNP in a certain position (not a gap).

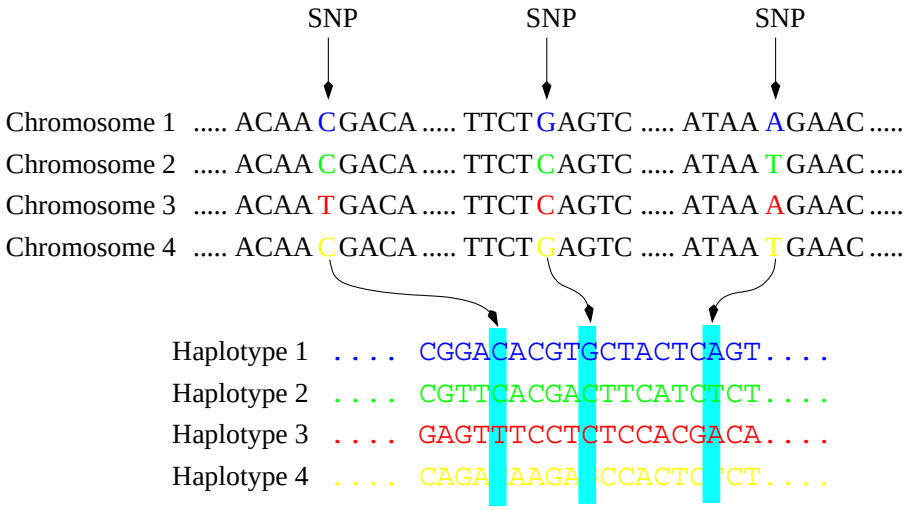

Figure 1: A family of four chromosomes in which the SNPs are highlighted is shown on top. Below the corresponding haplotypes.

Given the SNP matrix, we can define the conflict graph as follows: for each fragment there is a vertex; if two fragments are in collision, insert an edge between the corresponding vertices. When the SNP matrix contains no errors, conflicts can only arise between fragments originating from different haplotypes; therefore, the corresponding conflict graph is bipartite.

In the absence of errors or gaps, with sufficiently high coverage, this problem is easily solved. However, more realistic

models of the problem take into account several types and sources of error/noise in the data, and in this setting the problem is more challenging (intractable in some cases). Several algorithms and heuristics have been proposed in the literature to solve the haplotype assembly problem in this more realistic setting.

A common feature of these algorithms is that, due to lack of a common benchmark and of publicly available implementations in a unified framework, experimental comparisons between different methods are based on a narrow choice of data, parameters and algorithms. So far, in large scale haplotyping projects, insufficient attention has been given to the problem of how to identify the most suitable algorithm for a specific experiment and/or technology. To remedy this situation, and to provide a service to the bioinformatics community, we have developed the easy-to-use Web-based tool ReHap for testing diverse algorithms, on a variety of data sets, and with a variety of parameter settings.

Within ReHap our research group has developed a novel and fast heuristic method for the SIH problem called SpeedHap, which handles gaps efficiently, and is able to deal with high reading error rates (up to 20%) and low fragment coverage (as low as 3): situations in which other methods fail. We demonstrate these properties via experiments on real human data from the HapMap project. Away from the controlled environment of a lab it is likely that the current portable technology for sequencing will produce less reliable data. Moreover, if a real-time and high throughput response is needed to care for the needs of many individuals in a short time span, one might not be able to guarantee a high coverage of the fragments and low reading error rates. SpeedHap is a step towards the efficient extraction of useful information even from low quality data.

Rehap and SpeedHap have been developed by the Bioalgo Group at the Institute for Informatics and Telematics of CNR.

Link:
http://bioalgo.iit.cnr.it/rehap

Please contact:
Filippo Geraci
IIT-CNR, Italy
E-mail: filippo.geraci@iit.cnr.it

Sensor Network for Search and Rescue Operations in Collapsed Buildings

by Milt Statheropoulos

Earthquakes, technical failures and explosions are responsible for building collapses that result in victims becoming trapped. Second-Generation Locator for Urban Search and Rescue Operations (SGL for USaR, www.sgl-eu.org) is a mission-oriented project that aims at solving critical problems following massive destruction and large-scale structural collapses in urban locations. Urban search and rescue operations are carried out in a highly hazardous and unstable operating environment, with large numbers of injured, entrapped and vulnerable casualties needing help. In such situations, the first responders called to contribute to the search and rescue operations are usually outnumbered by the large number of collapses that they have to deal with in a short period of time.

Field operational tests at SDIS84 field terrain, France.

SGL for USaR follows a multidisciplinary approach in order to advance the state of the art in a wide spectrum of issues relating to urban search and rescue operations. These include effective and reliable casualty location methods; interoperability of data and systems; coordination between different units and disciplines; effective structural and environmental assessments of conditions in and under the ruins; medical monitoring of trapped casualties; and the location and differentiation of deceased casualties.

The main tasks of the project include: the development of simulation environments for testing location methods and devices; the development and validation of portable devices for search and rescue operations; the development and validation of smart sensors for monitoring conditions under the

ruins; the development of an ICT platform that will integrate the received data; the control of the information flow from the field to an operational command and control centre; and the management of medical information, including privacy and bioethics. A first responders' standalone portable device was designed during the first year of the project and will be developed and validated in the following years. This device will integrate optical, acoustic and chemical methods for locating entrapped victims and detecting hazardous conditions. The portable device will be able to run on-site (field) chemical analysis to detect chemical signatures both from victims and unsafe conditions. Field Chemical Analysis is employed for the first time in search and rescue operations for detecting signs of life as well as chemical signs from dead bodies.

The project will also develop a network of sensors to allow unattended monitoring of the conditions in the ruins of a collapsed building. It will integrate information collected from cameras, microphones and chemical sensors with the aim of detecting signs of life and hazardous conditions signatures in the collapsed building. We are also attempting to develop a distributed and heterogenous wireless sensor network. This network will employ ICT principles such as the Internet Protocol (IP) for long-distance communication among nodes, whereas RF-based technology, being lightweight and with low power requirements, will wirelessly connect short-distance nodes. An Integrated Platform will act as a command and control centre, responsible for managing data acquired by the devices of the network. The platform will integrate all the previous data, ensure interoperability and control the flow of the information from the field to the operational centre. Furthermore, by using special antennas it will provide reliable telecommunication facilities to the first responders.

In SGL for USaR data fusion algorithms will be applied according to the following levels:

1. Sensor fusion on a single device, where each device is equipped with different sensors. Sensors can work cooperatively or complementarily, and in both cases, data fusion brings up a more complete picture of the environment.
2. Sensor fusion over similar devices, where common devices are grouped to cover a larger range.
3. Sensor fusion over different devices, where the type of information is no longer similar. This is the most powerful way of combining sensor information to monitor the environment under the ruins.

This multilevel and multifaceted information fusion system will integrate a number of algorithms for vital sign detection, such as: adaptive threshold detection, adaptive motion detection, image and sound enhancement, image blob detection, alarm assessment, sound and video compression, etc. The corresponding context-dependent soft computing techniques attempt to take advantage of a dataset's strengths while minimizing the impact of its weaknesses. Its context-dependent multilevel algorithmic design is also based on the requirement for processing power, communication bandwidth, power supply and storage features. The general characteristics of this critical system will be interoperability, flexibility, reconfigurability, extensibility, modularity, evolvability and robustness, making it suited to a large field of applications.

For disseminating the results of the project, a Web page has been developed (http://www.sgl-eu.org) in which a technology forum is operating. The basic goal of the forum is to become a platform for the exchange of experience, knowledge and know-how within the SGL for USaR project.

SGL for USaR is an integrated and multidisciplinary project with 21 partners including rescue teams, researchers and SMEs. It is a four-year FP7 security project that started in October 2008. It is coordinated by the National Technical University of Athens, Greece (Field Analytical Chemistry and Technology Unit/FIACTU).

Link:
http://www.sgl-eu.org

Please contact:
Milt Statheropoulos
National Technical University of Athens, Greece
E-mail: stathero@chemeng.ntua.gr

IT Security: Risk-Based Prediction Tool and Method for Critical Infrastructures

by Jocelyn Aubert, Christophe Incoul and Djamel Khadraoui

In today's world, where most critical infrastructures are based on distributed and interdependent systems, security failures have become very common, even within large corporations. The FP7 MICIE (Tool for systemic risk analysis and secure mediation of data exchanged across linked CI information infrastructures) project aims to design and implement an alerting system that identifies, in real time, the level of possible threats induced on a given critical infrastructure (CI) by 'undesired' events occurring in the CI and/or other interdependent CIs. Here, we briefly describe the approach of a real-time risk-based monitoring solution for interdependent services of CIs.

Current risk analysis methods do not provide a way to share risk knowledge between providers who together form a CI. Providers have expertise on risks to their own infrastructure, but not on those to related infrastructures of other providers. In most cases, this knowledge cannot be shared between providers for confidentiality reasons. The idea of the approach developed by The Public Research Centre Henri Tudor, in the MICIE project, is to limit the exchange of data between CIs to the status of the security properties related to shared services. Confidential pieces of information about the CI (infrastructure, tools, equipment) are not exchanged. To reach these goals, we have developed a risk-based methodology that will monitor interdependent services based on generic risk and assurance levels using the classical security

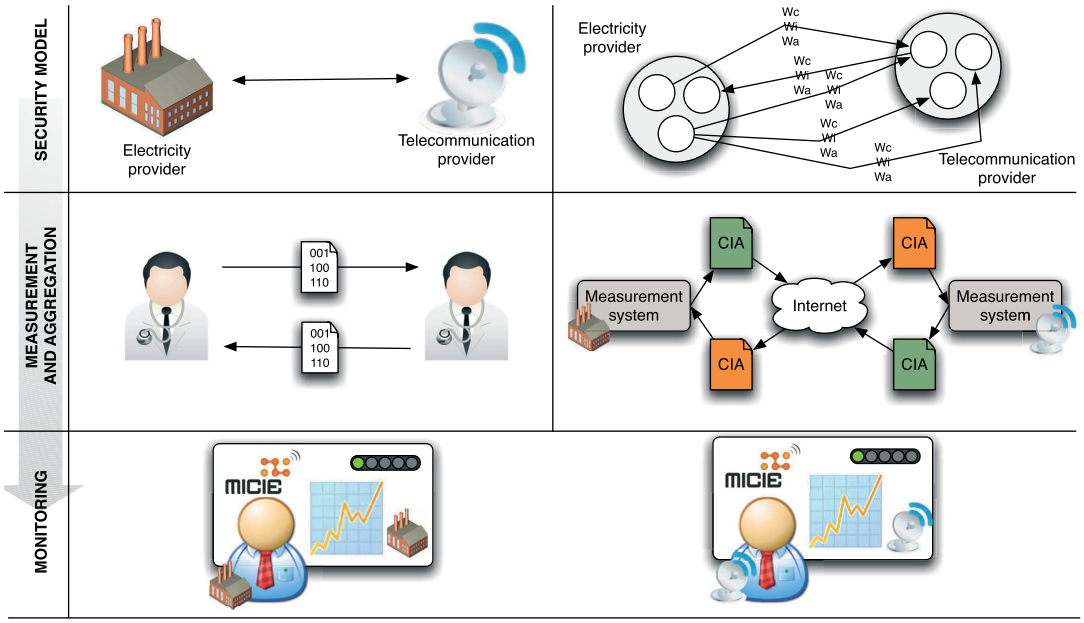

Figure 1: IT security risk-based prediction methodology for CI.

Some icons in Figure 1 were designed by DaPino Colada and were extracted from www.iconfinder.net under CC NY 3.0 license.

properties: confidentiality, integrity and availability. This allows each CI owner to monitor, react to and adopt the best behaviour corresponding to the security status of its different services.

A three steps methodology

The methodology allows a realistic representation of the security properties of interdependent systems. Using security properties to abstract the physical implementation will allow a wider range of systems to be addressed since the security objectives are the same for all systems. The security properties of an interdependent system are defined by the attributes confidentiality, integrity and availability. As illustrated in Figure 1, this methodology is composed of three steps: the security model step, the measurement and aggregation step and the monitoring step. These three steps are described below.

Security model

The whole methodology is based on an efficient and realistic modelling of the infrastructure paired with a global estimation of the dependencies between services. This service-oriented modelling aims to identify and represent all the critical services within the infrastructure as their links on external services in terms of confidentiality, integrity and availability. Using this model, the infrastructure will be able to define a generic level of risk and assurance of the different services that will be exchanged among interdependent CIs. Modelling relies on a reductionistic representation of the services, combining a functional model that defines and balances in terms of importance the relations between each critical service, with the results of a risk analysis made on each CI that aims to identify the main critical services. The reasoning based on the security model is only possible if all the CIs share a common representation of assurance and risk levels; for this, generic scales for both assurance and risk levels were defined.

Measurement and aggregation

Using the security model, the developed system is able to continually perform measurements on the various components composing monitored services. The results of such measurements are normalized and aggregated in the form of risk levels for each of the considered security attributes. Normalization and aggregation are realized thanks to advanced algorithms.

Monitoring

Using the security model, each CI sends local computed risk levels to each of the services depending on it. A service that receives a risk level can compute a risk linked to its dependencies, and thus update its global risk level by combining its own local data with the data provided by services on which it depends. This consolidated information enables the CI operator to integrate risks induced by remote services on the global CI risk estimation. Doing so, the operator is able to identify and to act and prevent security incidents from propagating in the interdependent critical services.

Experimentation

The methodology is the subject of experiments. An application prototype has been developed and is based on a multi-agent system, in charge of collecting, normalizing and aggregating all the considered measures using advanced algorithms. These measures are then combined in order to provide each CI owner with consistent consolidated information on the current level of risk.

The first results show that such an approach can provide a more comprehensive view of CIs and their interdependencies. However, enhancements could still be considered, as for example a definition of dynamic interdependency weights, to provide flexibility according to the current situation.

Links:
http://www.micie.eu
http://citi.tudor.lu

Please contact:
Djamel Khadraoui
CRP Henri Tudor, Luxembourg
Tel: +352 42 59 91 1
E-mail: djamel.khadraoui@tudor.lu

Formal Modelling and Analysis of Predictable Java

by Thomas Bøgholm, René R. Hansen, Anders P. Ravn, Hans Søndergaard, and Bent Thomsen

While embedded systems development is becoming increasingly complex due to demands for greater functionality and a shorter time to market, it is still using low-level, close to hardware, implementation languages. Modern languages like Java handle such complexities more elegantly, but issues with predictability hinder their adaptation in the embedded systems world. We have developed Predictable Java, a Java-based framework for safety-critical embedded systems, along with analysis tools based on formal modelling.

Embedded systems are everywhere, and are increasingly affecting our everyday life. Demands for shorter time to market, greater functionality and lower costs makes it difficult to develop these systems using traditional programming languages. Years ago, the traditional software industry shifted from low-level languages to modern languages with great success, increasing productivity significantly. Standard Java, however, is difficult to deploy and analyse in a real-time system setting. Java is based on a rather complex virtual machine, making analysis hard, and being object oriented, it has a very dynamic behaviour with almost unpredictable memory and time consumption. Additionally, Java is based on garbage collected memory, causing unforeseen interruptions which further complicates analyses.

We suggest that modern software engineering practices and languages be used in embedded systems software development by providing the required technology. The framework based on our Predictable Java profile provides the ability to express very complex systems in a simple, understandable and clear language which is recognizable and easily adaptable by most Java programmers.

Our tool SARTS (schedulability analsyser for real-time systems) is a schedulability analysis tool based on a formal model of timed automata, for proving schedulability of systems developed using our framework. This work is funded by DaNES (Danish Network for Intelligent Embedded Systems) and CISS (Center for Embedded Software Systems) at Aalborg University, and is publicly available.

In an attempt to reach the closest correspondence between actual running code and the properties being verified, SARTS works directly on the actual executed bytecode. This allows for a very close correspondence between the running system and the verified properties. For our prototype, we use the time-predictable Java processor, JOP, developed by Martin Schöberl at the Vienna University of Technology. This is an open and well-documented processor designed with predictability in mind.

The bytecode from the Java compiler is converted into a formal model, representing the real system, which is analysed in UPPAAL: a tool for modelling and verification of real-time systems.This approach provides a fully automatic process from code to verdict; developers need no training in formal verification, timed automata or UPPAAL.

Our technique provides a tight analysis by taking into consideration the control flow in the program code. This allows for analysing blocking, preemption and many of the dynamic features available in Java, which are considered hard to analyse, such as finalizers and dynamic method dispatch.

The Predictable Java framework provides a target for developing other analyses. We are designing the language profile as well as the analyses for the language profile. This enables us to fine-tune the developer framework in order to reach a reasonable compromise between restrictions and analyz-

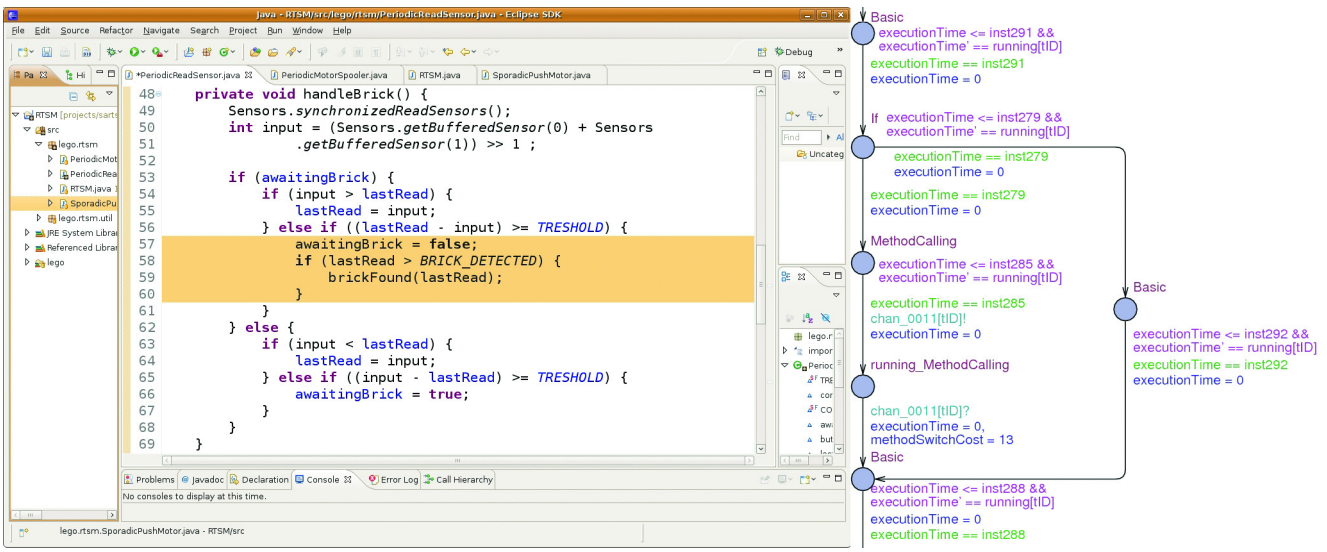

Figure 1: (left) A small snippet of a Java code verified using SARTS; (right) a simplified UPPAAL model of the highlighted code.

ability, making Java suited for safety-critical embedded systems.

Inspired by the techniques applied in SARTS, we are also developing memory consumption analysis, along with Eclipse IDE plugins that will assist programmers using the profile and profile conformance checker.

Links:
SARTS and Predictable Java: http://pj.cs.aau.dk
DaNES: http://danes.aau.dk
CISS: http://www.ciss.dk
UPPAAL tool: http://uppaal.com

Please contact:
Bent Thomsen
Aalborg University / DANAIM, Denmark
Tel: +45 9940 8897
E-mail: bt@cs.aau.dk

Hybrid Systems Design Methodology

by Michal Pluska and David Sinclair

Embedded systems that interact with the physical world should be designed with a high degree of safety. In most existing design approaches that deal with such systems, verification is done afterwards and may lead to redesign problems. Moreover verification is done by simulation which may not consider the whole spectrum of system operation. The Lero – Irish Software Engineering Research Centre work at the Dublin City University introduces a methodology that helps design systems that can be verified by formal methods and according to system requirements. Moreover it helps partition the system into the objects from which it is built.

Hybrid systems are built as a combination of computational systems and real-world physical parts. The computational parts can be seen as one or more embedded computers connected by a network and interacting with the physical world via sensors and actuators, with control engineering techniques helping to coordinate this set-up. The computational part of the system is therefore affected by the real-world physical part; moreover, it has to deal with the real-time properties of the physical world. In addition, hybrid systems can be seen as a group of cooperating devices of which the computing device is only one; on the other hand, it may be the one that is responsible for coordinating the whole system and ensuring it works as intended.

The complexity of hybrid systems makes the design a very challenging issue. Obtaining the correct design can be problematic, and extensive testing of various prototypes may be necessary. Unfortunately this can raise the design costs and the time required to finish it. The other problem to tackle during the design is correctness of the system. This problem is addressed with a design methodology using formal methods for verification of the system in the design.

Most widespread approaches to hybrid system design and later verification reassemble the bottom-up design methodology. The work begins with information about the system gathered in the form of physical rules and equations. Objects in the design are related to physical objects of the system and the interactions between them are made according to the physical equations. Those interactions can also be described as an agent-based approach to the description of a system. This bottom-up methodology is focused on detailed description and its relevance to the physical world.

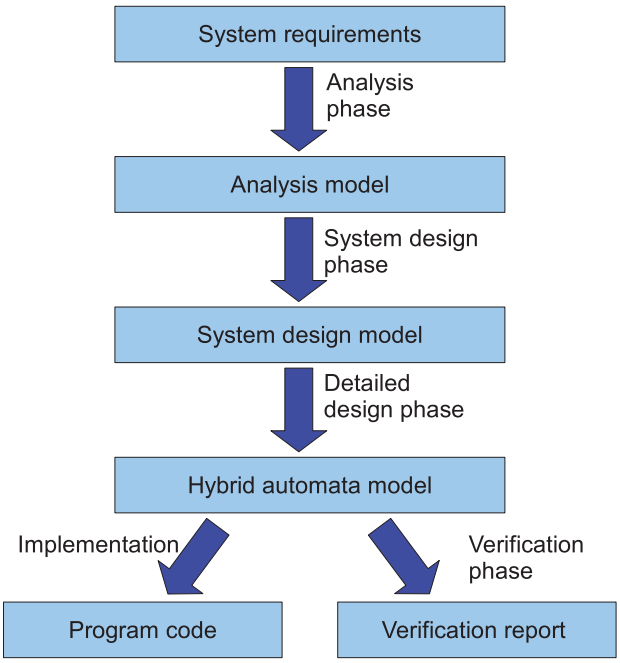

Figure 1: The proposed modelling and verification of the hybrid system.

Describing the system from the beginning of the design may lead to too many detailed equations in the initial stage. This is clearly visible with large systems described by many physical laws. Currently abstraction of those systems, useful during design, is possible only by using less realistic physical equations. Moreover there is no established rule for keeping track of the changes. Every case is explored individually with different levels of abstraction, and its solution is found by numerous experiments. The number of considered details can be problematic for the design engineer and can easily be a cause of error.

System design by this approach is verified by numerous simulations. The aim of simulation is to avoid extensive testing after the manufacture and hence reduce the cost and time. Unfortunately the degree of confidence in the correctness of the simulated design may be limited. The main cause of this is that too large a set of input data will cause unpredicted interactions with the environment that are impossible to

check. Building a prototype of the system suffers from identical problems as the system complexity rises.

In comparison, formally verifying high-level designs of complex systems can be very useful for hybrid systems, many of which are safety critical. By building a formal mathematical model of the system it is possible to use automated model-checking methods to prove that all requirements are met or all possible system input sequences are checked. Simulations only allow some of the potential system inputs to be checked.

On the other hand, the usefulness of formal methods is limited by the lack of a well-defined methodology that would make it broadly applicable. There are no tools for interactive model building and analysis interpretation. Moreover the complexity of the system in design can be overcome by using appropriate abstraction of model's details what is currently not supported at all. There is also a need for aids to translate informal requirements specifications into formal specifications. A final aspect worth mentioning is that formal methods specification and verification might be problematic for practical engineers, because their focus during design is on different aspects.

The proposed modelling and verification of the hybrid system is divided into stages (see Figure 1). This partitioning allows the correct complex system to be built from the start and avoids problems which otherwise would only be found after the initial design is complete. Our design methodology is focused on describing the system requirements by examples of its usage in use cases. It allows verification of the gathered requirements and represents a starting point for the analysis. The analysis of those objects will hierarchically decompose them according to the abstraction levels that build the system model. This is done with respect to the abstraction level boundaries. This stage of the methodology design system model has all the features described in the requirements and is a backbone of the model used in the system verification process. All the necessary information is gathered during the design of the system model. The final verification of the system and the possible parameters of the system is done by a hybrid automaton. This approach builds and verifies the system model independently of the hardware. This makes it possible to use control software already verified on a different hardware platform, depending on the needs of engineers.

This work was supported, in part, by Science Foundation Ireland grant 03/CE2/I303_1 to Lero - the Irish Software Engineering Research Centre (http://www.lero.ie)

Please contact:
Michal Pluska
Dublin City University, Ireland
E-mail: michal.pluska2@mail.dcu.ie

Content-Sharing for Mobile Social Networks in the Haggle Platform

by Marco Conti, Franca Delmastro, Andrea Passarella

In mobile pervasive environments, mainly characterized by intermittent connectivity conditions, gathering social and context information about the users and their behaviour is fundamental to providing efficient communications and content exchange.

The emerging trend in social network applications and active participation of users in generating and sharing content (as defined by the Web 2.0 paradigm) further increases the need for users to be continuously connected with each other and to the Internet, everywhere and always. However, users' mobility and the physical limitations of wireless communications do not allow us to apply the same principles and solutions developed for the conventional (static) Internet. We cannot assume a single and fully connected network where mobile devices can communicate through a stable path. For this reason, in the last few years, Opportunistic Networking has emerged as a new communication paradigm to cope with this kind of scenario and especially with intermittent connectivity conditions.

The European FET-SAC (Future and Emerging Technologies - Situated and Autonomic Communications) Haggle project, started in 2006 and currently reaching its end, has been one of the pioneering investigators of this topic. It defines a new autonomic networking architecture to support communications and content exchange between users and mobile devices that could never be connected through a conventional multi-hop path, due to the users' mobility or significant partitions of the network. The main concept is to exploit users' mobility and their patterns of interaction to generate routes on the fly. Specifically, every time a mobile device receives a message with a specific destination, it selects from among its current 1-hop neighbours the best forwarder by analysing several elements (eg the probability of arriving at the destination in the near future, the common interests between the local user and the destination etc). Thus, the system processes additional information to complement knowledge on the network topology that is intrinsically unstable. Specifically, the best actions to be taken by networking protocols when two users are in contact are identified through knowledge of their social behaviour, and of context information related to (i) the user (eg habits, interests, timetables); (ii) the mobile devices (eg capacity, battery lifetime); and (iii) the services running on them (eg list of shared files, set of interesting genres in case of content sharing). For this reason social and context information assumes a fundamental role both in the network protocols, in order to guarantee efficient communications among users, and in the design of efficient services and applications tailored to mobile pervasive environments.

The Social-Aware Content Sharing service we designed and developed in the framework of Haggle exploits a context

definition tailored to both the features of the specific service and those of opportunistic networks. The main idea is that users participating in the service can decide to declare to other participants information on the content they want to share, plus some personal information that can help the system predict their social interactions and consequently mobility patterns. The user directly configures this personal profile through the graphical interface shown in Figure 1.

We presume that users are grouped in social communities characterized by some common interests (eg coworkers, flatmates, family, friends). Each user can belong to one or more communities, and users of the same community have strong social links with each other, periodically getting in touch. This information is used by the service to identify possible interesting content, even for users who are not physically connected, and more generally by the system to improve forwarding decisions on the network by probabilistically predicting future contacts among users.

As a practical example of our service we can consider the scenario depicted in Figure 2. There are two communities (X and Y) whose users have some common interests in sharing files (eg users A, B and C are all interested in mp3 files, while D and E are interested in avi and jpg). Every time a mobile device of one of these users comes into contact with a new neighbour, they exchange context information (mainly related to user and service data) so that each device knows the interests and the shared files of all its neighbours. Thus, users are able to request from and exchange content with the device of their current neighbours. This information is maintained on the device as a context history, even when the user moves to another community. For example, when user C moves to community Y and their mobile device gets in touch with the new neighbours, it exchanges context information and analyses whether some of the available content might be useful to previously encountered users. Since it is likely that it will return to the original community, the device autonomously decides to fetch some of this content (depending on its utility for the other users and on the current status of its resources), so that when it returns, this content

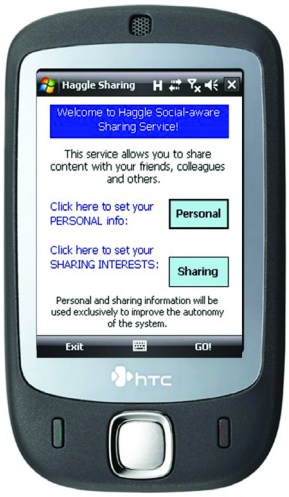

Figure 1: Interface of the Social-Aware Content-Sharing Service.

can be shared with users of community X. In this way, users who have never been in direct contact can exchange interesting content based on their context information.

Beyond the specific example of content sharing, social-aware context information is integrated in the Haggle platform in order to allow all the network components to exploit it. We integrated a Context Manager in the Haggle architecture in order to acquire and manage context data. As shown in Figure 3, the Haggle architecture consists of a set of managers implementing the main features needed to communicate and deliver messages on the network, and a central kernel through which they can interact to improve efficiency, thus providing a completely layer-less architecture above the data link. The Context Manager is able to manage context information and make it available both to other interested managers and upper-layer services and applications in order to improve their features in mobile pervasive environments.

Link:
EC Haggle project: http://www.haggleproject.org

Please contact:
Marco Conti, Franca Delmastro or Andrea Passarella
IIT-CNR, Italy
E-mail: firstname.lastname@iit.cnr.it

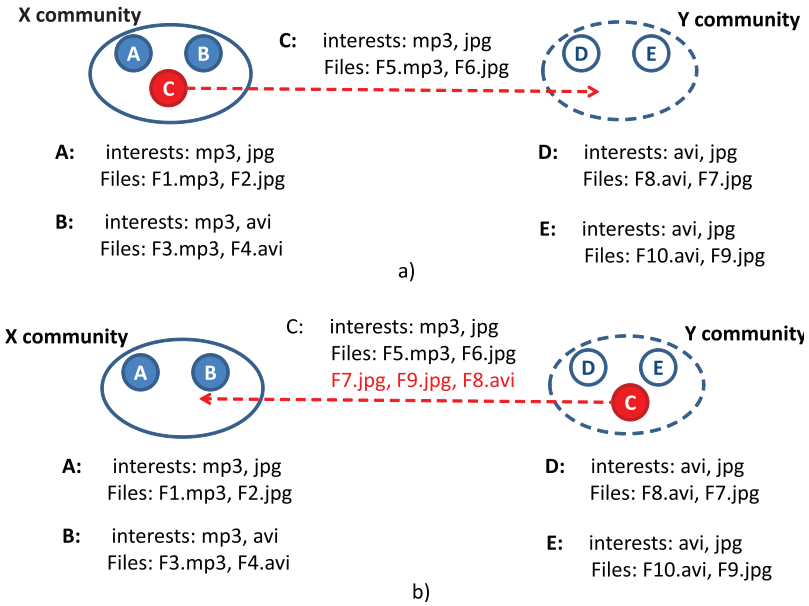

Figure 2: A realistic scenario.

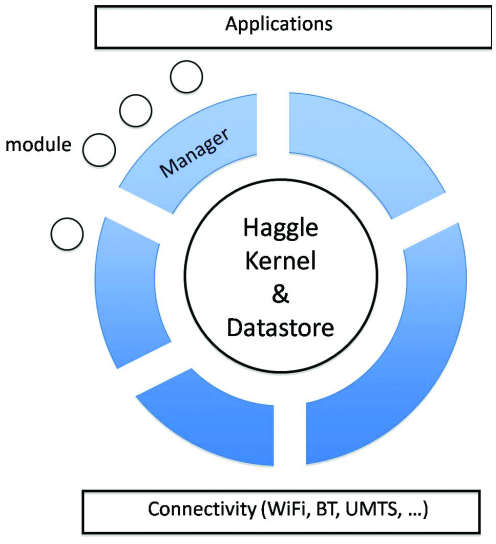

Figure 3: The Haggle architecture.

EXTRA Helps SMEs Assess their Knowledge Management Practices

by Sanae Saadaoui and Frederic F. Monfils

EXTRA has been designed to help small and medium enterprises (SMEs) improve the way they develop and deliver software by making better use of each company's previous experiences. In order to do this, SMEs must record these experiences. While knowledge management evaluation used to be inaccessible to SMEs, this is no longer the case. EXTRA has developed a self-assessment tool providing a set of good practices in software development for SMEs.

Among the various initiatives envisioned so far to improve the quality of software products, knowledge sharing has been under recent investigation. The goal is to record experiences from past projects, both successes and failures. Capitalizing on experiences accumulated by the company can significantly improve the quality of the product, the efficiency of the team and the satisfaction of the customer. However, existing solutions in knowledge management are expensive and often require large investments of time, resources and money. This explains why these solutions are seldom applicable in SMEs.

The EXTRA project aims at producing a set of knowledge management techniques tailored to SMEs that develop and/or deliver software, hence improving the quality of their services. These techniques are based on theoretical research in this area. EXTRA follows the 'Action Research' methodology. In this methodology, the researchers and the participating SMEs collaborate to develop and test innovations in a local context. This process ensures the commitment of players and the quality and relevance of research activities.

The EXTRA project will finish in a few months, and a set of techniques have now been validated. Following is a description of the main achievements which led to this validation and to the self-assessment tool developed during the project.

The first step consisted in surveying the requirements of the SMEs participating in the project as well as understanding their current experience with knowledge management. A questionnaire was designed to determine current and future interests in five major categories of strategies for knowledge management:

- the System school, which uses information systems (intranet, wiki) to codify knowledge
- the Cartographic school, which records employees and their skills
- the Engineering school, which records how to perform the activities and which particular role performs each task (this is similar to the writing of procedures when initiating an ISO 9001 certification)
- the Organizational school, which uses networking and meetings to share experience on a common topic
- the Spatial school, which arranges the work environment to encourage and promote exchange among employees.

The results showed a strong preference for the 'System' and 'Engineering' schools, with interest in the 'Spatial' school being low.

Based on the results of the survey, a first version of the guide was created that covered seven techniques. The descriptions

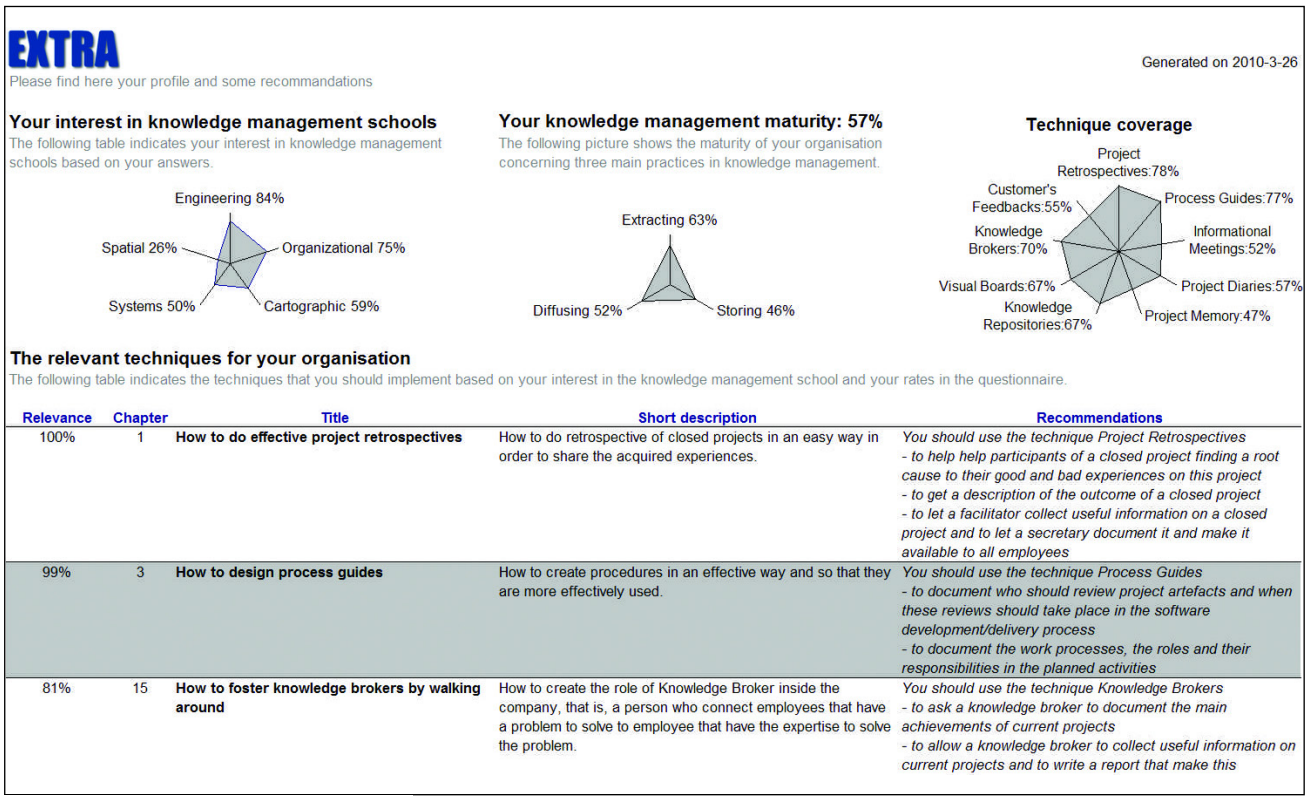

EXTRA

Please find here your profile and some recommandations

Generated on 2010-3-26

Your interest in knowledge management schools

The following table indicates your interest in knowledge management schools based on your answers.

Your knowledge management maturity: 57%

The following picture shows the maturity of your organisation concerning three main practices in knowledge management.

Technique coverage

The relevant techniques for your organisation

The following table indicates the techniques that you should implement based on your interest in the knowledge management school and your rates in the questionnaire.

Relevance	Chapter	Title	Short description	Recommendations
100%	1	**How to do effective project retrospectives**	How to do retrospective of closed projects in an easy way in order to share the acquired experiences.	*You should use the technique Project Retrospectives - to help help participants of a closed project finding a root cause to their good and bad experiences on this project - to get a description of the outcome of a closed project - to let a facilitator collect useful information on a closed project and to let a secretary document it and make it available to all employees*
99%	3	**How to design process guides**	How to create procedures in an effective way and so that they are more effectively used.	*You should use the technique Process Guides - to document who should review project artefacts and when these reviews should take place in the software development/delivery process - to document the work processes, the roles and their responsibilities in the planned activities*
81%	15	**How to foster knowledge brokers by walking around**	How to create the role of Knowledge Broker inside the company, that is, a person who connect employees that have a problem to solve to employee that have the expertise to solve the problem.	*You should use the technique Knowledge Brokers - to ask a knowledge broker to document the main achievements of current projects - to allow a knowledge broker to collect useful information on current projects and to write a report that make this*

Figure 1: Self-assessment tool.

of the techniques were validated by the participating SMEs, and some were implemented by them. The second version of the handbook incorporated the feedback of these implementations.

It appeared during the various meetings with SMEs, that there was a high demand for a self-assessment tool to help them evaluate their current practices in knowledge management. CETIC was appointed to develop this assessment tool given its proven expertise in assessment methods and software process improvement targeted to small companies. CETIC also actively participates in the ISO WG24 working group and this group is responsible for developing a new standard, ISO/IEC29110, suited to small and very small businesses.

The self-assessment tool consists of a questionnaire. The benefits of this assessment are threefold. First, it identifies the interests of the company in the five main knowledge management schools; second, it evaluates the maturity of the company in the three main practices (capitalization, recording and diffusion of knowledge); and finally, it directs the company not only to the most appropriate knowledge techniques described in the handbook but also to general recommendations in knowledge management (see Figure 1).

This tool is currently available as a spreadsheet and should be soon made available as an online Web form with direct guidance.

Final presentation of the EXTRA project results

The results of the EXTRA project will be presented at the XP2010 conference to take place in June 2010 in Trondheim (Norway). A session will be dedicated to the project. There will also be presentations and workshops held in order to guide SMEs in their implementation of knowledge management techniques and in the evaluation of their IT-related practices.

Acknowledgement

EXTRA is a European CORNET project funded by DGTRE - Walloon Regional Ministry – Directorate General for Technologies, Research and Energy, Belgium. It involves partners from three different countries: Belgium (CETIC), Norway (ICT Norway, SINTEF) and Cyprus (CITEA, Inteliscape, Virtual IT). The project started in May 2008 and runs for a period of 24 months. In addition to research partners, participating SMEs based in Cyprus, Norway and Belgium's Wallonia area are also involved in order to validate the research.

Links:

http://www.cornet-extra.eu
http:// www.cetic.be

Please contact:

Sanae Saadaoui, CETIC, Belgium
Tel: +32 71 490 748
E-mail: sanae.saadaoui@cetic.be

Frederic F. Monfils, CETIC, Belgium
Tel: +32 71 490 736
E-mail: frederic.fleurialmonfils@cetic.be

New Headway for the Semantic Web

by Stephanie Parker

"I have a dream for the Web in which computers become capable of analysing all the data on the Web", said Tim Berners-Lee in 1999. Therein lay the future. Linked Data is part of the vision of the semantic web to enrich the structure of the Web by embedding semantic annotations into data to improve the quality of search, collaboration, publishing and advertising and enable applications to become more integrated and intelligent.

Initiatives recently spearheaded by the UK and U.S. governments to make non-personal public data available online are aimed at increasing accountability and raising awareness of government functions among citizens. Linked Data is a way of making that information available easily and efficiently. In general, it is open, modular and scalable. Because linked data is expressed in open, non-proprietary formats it is accessible through an unlimited variety of applications. This can also be combined (mashed-up) with other pieces of linked data, with the option of adding more to existing data even when the terms and definitions used change over time.

The Resource Description Framework (RDF), a W3C standard for describing resources in the Web and a major component in what is proposed by the W3C for the Semantic Web, serves as the "interconnection bus" for current data formats. One of the distinguishing features of the linked data technology is that it allows data communication to be composed of mixed vocabularies, which come from a community, be it local, state, national or international. RDF uses Uniform Resource Identifiers (URIs) to allow the linking of things and concepts. The URI thus helps identify a resource on the internet, enabling interaction with representations of the resource over a network using specific protocols. RDF also allows interoperability to be added wherever this is cost-effective.

"The Semantic Web will transform the World Wide Web into a more useful and powerful information source. In particular it will revolutionise scientific and other web publishing by defining new web technologies that make more web content accessible to machines. These technologies will provide better tools that make it easier for people to create machine-readable content that is widely available", said Professor Dame Wendy Hall, University of Southampton.

Digital Libraries can benefit from RDF to facilitate digital resource management and support knowledge management for an interoperable information environment like that found in a digital library (DL). A good case in point is the Digital Library of the Association for Computing Machinery (ACM), the world's largest, most comprehensive and respected on-line resource in the field of computing, which is due to adopt RDF in 2010. Providing the full text of every article from every journal, magazine, conference proceedings, newsletter, oral history interview and video published since 1954, the ACM Digital Library has been at the fore-

front of technology since its inception. “The introduction of Semantic Web tools in the ACM Digital Library will enable its hundreds of thousands of professional and student users to more easily find, share, and combine information on the Web", explained Dame Wendy Hall, President of ACM.

The potential for RDF is huge. Media and publishing are among the early adopter sectors with use cases from the BBC and The New York Times, which refers to its Linked Data as its “Treasure Map”. Life science publishers have also adopted this model and cutting-edge research, information on climate change and academic publishers are all beginning the transition. All these early use cases are demonstrating new ways to manage and discover information. The ultimate goal is to make semantic web mainstream and enable information professionals through dedicated courses. Some of the challenges involve the security architecture of data, which can be best addressed by learning from lessons in the past.

“Linked data could be an even bigger sea change than the world wide web not only because it plays a key role in assisting a wide range of people in government and research, but also socially”, said Dame Wendy Hall. With Linked Data it will be much easier to tackle grand global challenges like climate change, energy and health issues, ageing, and world poverty by sharing data to capture correlations and trends more effectively and more quickly, spotting clues that enable researchers to make that all-important step forward. Dame Wendy Hall remarked, “Linked Data means lots more innovation but attitude is as important as the technology to get critical mass. The key thing is to arrive at universal standards just as we did with the web”.

The Semantic Web is not a separate web but an extension of the current one. The UK government’s recently announced plans to create a new institute for Web Science with Tim Berners-Lee at its helm and hosted by Oxford and Southampton universities, means that the vision is now a step closer. Imagination is our only barrier.

Links:
ACM: http://www.acm.org/
ACM: Digital Library http://portal.acm.org/dl.cfm

Online Information Conference, 1-3 December 2009, London – keynote by Wendy Hall and Nigel Shadbolt, University of Southampton. W3C – RDF and on-line government data Podcast at http://2009.online-information.co.uk/online09/podcasts.html

Please contact:
Stephanie Parker
Trust-IT Services
E-mail: s.parker@trust-itservices.com

Event Report

36th SOFSEM: Current Trends in Theory and Practice of Computer Science

by Julius Stuller

SOFSEM is a successful established annual international winter conference devoted to the theory and practice of computer science. Its aim is to present the latest research developments and so foster cooperation between professionals from academia and industry in various areas of computer science. SOFSEM 2010 was held in Hotel Bedøichov, Špindlerùv Mlýn, Czech Republic, on 23-29 January 2010.

SOFSEM (historically from SOFtware SEMinar) is the foremost Czecho-Slovak computer science conference. It is organized every third year in Slovakia, with the intermediate years being held in the Czech Republic.

SOFSEM offers a unique opportunity to quickly obtain a representative overview of the areas that are selected as the topics of the year. By tradition, the program consists of a dozen invited talks by prominent researchers, contributed talks selected by the program committee and the Student Research Forum. SOFSEM is organized in plenary and parallel tracks, with one track being devoted to the foundations of computer science. SOFSEM 2010 offered in addition three outstanding tracks: Principles of Software Construction; Data, Knowledge and Intelligent Systems; and Web Science. The schedule is tailored so that there is the opportunity to interact with other participants (invited speakers, researchers, students).

SOFSEM conferences are well known for the high quality of their invited speakers. In this year’s edition, eleven invited speakers gave very interesting talks:

Foundations of Computer Science
- Parosh Aziz Abdulla, Uppsala, Sweden: Forcing Monotonicity in Parameterized Verification: From Multisets to Words
- Yishay Mansour, Tel Aviv, Israel: Regret Minimization and Job Scheduling.
- Roger Wattenhofer (Zurich, Switzerland): Theory Meets Practice, It's about Time!

Principles of Software Construction
- Arie van Deursen (Delft, The Netherlands): Research Issues in the Automated Testing of Ajax Applications.
- Oscar Nierstrasz (Bern, Switzerland): Lessons in Software Evolution Learned by Listening to Smalltalk.
- Andy Schürr (Darmstadt, Germany): Model-Driven Software Product Line Testing: An Integrated Approach.

Data, Knowledge, and Intelligent Systems
- Theo Härder (Kaiserslautern, Germany): Essential Performance Drivers in Native XML DBMSs.
- Ondøej Ku□elka and Filip □elezný (Prague, Czech Republic): Taming the Complexity of Inductive Logic Programming.
- Yannis Manolopoulos (Thessaloniki, Greece): Continuous Processing of Preference Queries in Data Streams.

Web Science
- Guus Schreiber (Amsterdam, The Netherlands): Web Science: the Digital-Heritage Case.

Jiri Wiedermann at the SOFSEM conference.

- Dave Raggett (Sophia-Antipolis, France/UK): The Web of Things: Extending the Web into the Real World.

Since 1995, the SOFSEM conference proceedings containing the invited and contributed papers have been published in the series Lecture Notes in Computer Science by Springer Verlag; this year's proceedings are the LNCS volume 5901.

SOFSEM is the ideal conference for discussions, for establishing personal contacts with colleagues, and for exchanging ideas. SOFSEM is well known for its familiar and inspiring atmosphere and is especially suited to young computer scientists.

An integral part of SOFSEM 2010 was again the Student Research Forum, organized with the aim of publishing and discussing students' projects. The forum offered students the opportunity to receive feedback on both the originality of their scientific work results and the work in progress. In total, SOFSEM 2010 was attended by some 120 participants, of whom more than 25 were students. Almost 80% of participants were not local.

As usual, SOFSEM also served as a platform for spreading ERCIM ideas directly. This is done via the presence of ERCIM President Keith G. Jeffery, his very active involvement in the SOFSEM Steering Committee, and through a small exhibition presenting ERCIM's activities.

SOFSEM 2010 would hardly be so successful without the enthusiasm, time, energy and all the activities of two persons: the SOFSEM 2010 General Chair Prof. Jan van Leeuwen and the SOFSEM 2010 Organizing Committee Chair Dr. Roman Spanek.

All those interested are cordially invited to the 37th conference in the series, SOFSEM 2011, which will be held in Slovakia, Hight Tatras, from January 22-28, 2011.

Links:
http://www.sofsem.sk
http://www.sofsem.cz

Please contact:
Julius Stuller
SOFSEM Steering Committee Chair
Academy of Sciences of the Czech Republic/CRCIM
E-mail: stuller@cs.cas.cz

W3C Offices Events

Internet New Year Event 2010 in Amsterdam

On January 14th the W3C Benelux Office hosted by CWI and many local internet and ICT-related organizations jointly organized a new years event in Amsterdam to celebrate the start of the new year. The event took place at NEMO (a Science Museum), Oosterdok 2 in Amsterdam. After a buffet at 18:00 the event started at 19:00 accompanied by a lot of parallel talks and activities. While the main space was reserved for drinks and conversation only there were a lot of interesting 'lightning talks' in the lecture hall. A highlight was the "Chairman's debate" with Olaf Kolkman (NLnet Labs), Erik Huizer (IPv6 Forum), Jan Willem Broekema (ISOC), Peter-Paul Koch (Fronteers), Cees de Laat (Gridforum), Rob Blokzijl (RIPE) and Bert Bakker (OpenDoc Society).

WGAG 2.0 Seminar

Already in December last year the W3C Benelux Office organized together with the Foundation Accessibility and ISOC.nl a seminar on Web Content Accessibility Guidelines (WCAG) 2.0. The event took place in the Royal Library The Hague. Michiel Leenaars (Director ISOC.nl and member of the board of Accessibility) chaired the event. Speakers were Eric Velleman (Foundation Accessibility), Paul Timmers (Head of Unit for ICT for Inclusion in the European Commission), Imke Arts - Vrijling (Ministery BZK), Raph de Rooij (ICTU), Robert Jan Verkade (EEND) and Shadi Abou Zahra (W3C). 73 people attended the seminar. More than half of them were from the government.

W3C Spain Office held three Workshops on W3C Web Standards

The W3C Spain Office organized three workshops on W3C Web Standards in November and December 2009. These workshops were held in three public Spanish universities: Universidad de Murcia, Universidad de Granada, and Universidad de La Laguna. All of them showed great interest in co-organizing the event, and they suggested important topics not enough represented in the current curricula such as Web standards or the W3C's role in the success of the Web.

ERCIM is the European host of W3C. Five of the nine European W3C offices are hosted by ERCIM members institutes.

More information:
http://www.w3.org/Consortium/Offices/

Calls for Participation

Multilingual and Multimodal Information Access Evaluation: CLEF2010 Conference and Labs

Padua, Italy, 20-23 September 2010

CLEF2010 aims at advancing the evaluation of complex multimodal and multilingual information systems in order to support individuals, organizations, and communities who design, develop, employ, and improve such systems. It is the continuation of the popular CLEF campaigns that have run for the past ten years and consists of two main parts: a peer-reviewed conference on experimental evaluation, which will innovate the CLEF tradition, and a series of Evaluation Labs, which will continue the CLEF tradition of community-based evaluation.

CLEF2010 Evaluation Labs

Experimental evaluation - both laboratory and interactive - is a key to fostering the development of multilingual and multimodal information systems that address increasingly complex information needs. The CLEF2010 labs propose a set of benchmarking activities and will complement these with workshops on emerging issues in evaluation methodology.

Benchmarking activities

- CLEF-IP: A benchmarking activity on intellectual property.
- ImageCLEF: A benchmarking activity on image retrieval.
- PAN: A benchmarking activity on plagiarism detection.
- RespubliQA: A benchmarking activity on question answering using multilingual political data.
- WePS: A benchmarking activity on web people search.

Workshops

- CriES: A workshop aimed at exploring the evaluation of searching for expertise in social media.
- LogCLEF: A workshop aimed at exploring methodologies for studying search engine log files.

Registration is now open via the CLEF2010 website.

CLEF2010 Conference

The CLEF2010 conference invites submissions on all aspects of multilingual and multimodal information access evaluation.

Relevant topics include, but are not limited to:

- Novel methodologies for the design of evaluation tasks, especially user-centric ones;
- Analysis of the impact of multilingual/multicultural/multimodal differences in interface and search design;
- Assessing multilinguality and multimodality in relevant application communities, e.g. digital libraries, intellectual property, medical, music, video, and social media.
- Alternative methods for improving and automating ground-truth creation, for example crowd-sourcing or clicklog-based;
- Prediction of success and satisfaction rates;
- Task-oriented metrics of success and failure;
- Evaluation of technology vs testing of scientific theories;
- Innovative and easy to communicate techniques for analysing the experimental results, including statistical analyses, data mining, and information visualization;
- Alternatives for and comparison of item-based, list-based, set-based, and session-based evaluation;
- Simulation (of queries, sessions, users) and information retrieval;
- Infrastructures for bringing automation and collaboration in the evaluation process;
- Component-based evaluation approaches;
- Evaluation and analysis using private or anonymized test data;
- Living laboratories and evaluating live systems.

Authors are invited to submit electronically original papers, which have not been published and are not under consideration elsewhere. Two types of papers are solicited:

- long papers: 12 pages max;
- short papers: 6 pages max.

Papers will be peer-reviewed. Selection will be based on originality, clarity, and technical quality.

The conference proceedings will be published in the Springer Lecture Notes in Computer Science (LNCS) series.
For information on submission details, see the CLEF 2010 website

Important Dates

- Submission Deadline: 2 May 2010
- Notification of Acceptance: 11 June 2010
- Camera Ready: 25 June 2010
- Conference: 20-21 September 2010
- Lab Workshops: 22-23 September 2010.

More information:
http://www.clef2010.org/

Call for Participation:

Evaluation Campaign on Plagiarism Detection and Wikipedia Vandalism Detection

held in conjunction with the CLEF'10 conference

Plagiarism detection in text documents is a challenging retrieval task: today's detection systems are faced with intricate situations, such as obfuscated plagiarism or plagiarism within and across languages. Moreover, the source of a plagiarism case may be hidden in a large collection of documents, or it may not be available at all. Informally, the respective CLEF-Lab task can be described as follows: Given a set of suspicious documents and a set of source documents, the task is to find all plagiarized sections in the suspicious documents and, if available, the corresponding source sections.

Following the success of the 2009 campaign and based on our experience we will provide a revised evaluation corpus consisting of artificial and simulated plagiarism.

Vandalism has always been one of Wikipedia's biggest problems. However,

the detection of vandalism is done mostly manually by volunteers, and research on automatic vandalism detection is still in its infancy. Hence, solutions are to be developed which aid Wikipedians in their efforts. Informally, the respective CLEF-Lab task can be described as follows: Given a set of edits on Wikipedia articles, the task is to identify all edits which are vandalism, ie, all edits whose editors had bad intentions.

Participants are invited to submit results for one or both of the tasks.

Important Dates:

- 1 March 2010: Training corpora release (Preliminary training corpora are alread available!)
- 3 May 2010: Test corpora release
- 1 June 2010: Result submission deadline
- 15 June: 2010:Notification of performance
- 15 July: 2010:Paper submission deadline
- 2 August 2010: Notification of reviews
- 1 September 2010: Final paper deadline
- 20-23 September 2010: Evaluation lab at CLEF conference.

More information: http://pan.webis.de

sponsored by ERCIM

SAFECOMP 2010

Vienna, 14-17 September 2010

The emphasis of the 29th International Conference on Computer Safety, Reliability and Security is to take a look at critical embedded systems which are already almost omnipresent, communicating and cooperating, interacting with each other, humans and environment. The topics include purely scientific/technical as well as societal aspects. It concerns (explicitly used or hidden) devices, services, infrastructures, machines, transport systems, robots etc. The conference will cover research in all related areas, industrial experience and practice, reports from interesting R&D projects, studies, standardization issues, assessment, evaluation and certification of systems, tools and utilities.

More information:
http://www.ocg.at/safecomp2010/

Call for Participation

CMPD 3 - Conference on Computational and Mathematical Population Dynamics

Bordeaux, France, 31 May - 4 June 2010

The Third Conference on Computational and Mathematical Population Dynamics (CMPD3) is the third joint meeting of the Conference on Mathematical Population Dynamics (MPD) and the Conference on Deterministic and Stochastic Models for Biological Interactions (DeStoBio), with a 25-year history of international meetings.

The aim of the meeting is to bring together people from different fields (applied mathematicians, computer scientists, biologists, clinicians, epidemiologists, ecologists, etc.) interested in (deterministic and/or stochastic) models for population dynamics and interactions.

Population dynamics here is intended in a very wide sense, including everything from populations of animals and plants, to populations of cells or molecules. Hence, topics of the conference will include ecology (including epidemic spread), cell population dynamics (including immunology, tumor growth, neurosciences) and molecular biology (including molecular evolution and genetics, genomics). The meeting will focus on modelling of quantitative data in these fields, analysis of models, and their applications.

Plenary speakers:

- Andre De Roos, Amsterdam University, Netherlands
- Arnaud Ducrot, Bordeaux 2 University, France
- Gabriela Gomes, Gulbenkian de Ciencia Institute, Portugal
- Yoh Iwasa, Kyushu University, Japan
- Mark Lewis, Alberta University, Canada
- Maia Martcheva, University of Florida, USA
- Sylvie Meleard, Ecole Polytechnique, France
- Andrea Pugliese, Trente University, Italia
- Mick Roberts, Massey University, New Zealand
- Sanyi Tang, Shaanxi Normal University, China

Invited sessions:

- Modern Developments in Mathematics of Infectious Diseases (Fabio Milner and Maia Martcheva)
- Biomathematics with emphasis on ecoepidemiology and spatio-temporal pattern formation (Ezio Venturino)
- Physiologically Structured Population Models (Odo Diekmann)
- Nosocomial transmission of infections (Christopher Kribs Zaleta)
- Matrix Models for Structured Population Dynamics (Dmitrii O. Logofet, Sandrine Charles)
- Stochastic population dynamics and birth-death processes (Yoram Louzoun)
- Mathematical Models for Emerging Diseases and Emerging Countries (Ying-Hen Hsieh and Hector de Arazoza)
- Noise-induced effects in population dynamics (Sergei Petrovskii, Horst Malchow & Michael Sieber)
- Dynamics of structures in cell biology (Stephanie Portets)
- Modelisation, Bioeconomy and fisheries management (Rachid Mchich, Nadia Raïssi and Pierre Auger)
- Epidemic models with household or network structure (Andrea Pugliese)
- From spatially explicit population models to mean-field dynamics (Andrew Morozov, Jean Christophe Poggiale)
- Modelling in immunology (Rodolphe Thiébaut)
- Modelling complex multi-strain dynamics in single and meta-populations. Application to influenza and other diseases (Sébastien Ballesteros, Bernard Cazelles and Elisabeta Vergu)
- Stochastic and Deterministic Population Processes: From Branching to the Transport Equation and Beyond (Marek Kimmel)
- Structured cells population dynamics models and related topics (Jean Clairambault, Pierre Magal, Glenn Webb)
- Modeling of Cancer Growth and Treatment (Evans Afenya)

More information:
http://www.sm.u-bordeaux2.fr/CMPD3/

Call for Participation

FMICS 2010 - 15th International ERCIM Workshop on Formal Methods for Industrial Critical Systems

Antwerp, Belgium, 20-21 September 2010, colocated with ASE 2010

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

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

The workshop is co-chaired by Marco Roveri and Stefan Kowalewski.

More information:
http://es.fbk.eu/events/fmics2010

Call For Papers

IWPSE-EVOL'2010 Joint ERCIM Workshop on Software Evolution and International Workshop on Principles of Software Evolution

Antwerp, Belgium, 20-21 September 2010

The IWPSE-EVOL workshop is the merger of the annual ERCIM Workshop on Software Evolution (EVOL) and the International Workshop on Principles of Software Evolution (IWPSE). The rationale for a common event is to capitalize on the synergies to be found when theorists and practitioners meet.

Research in software evolution and evolvability has been thriving in the past years, with a constant stream of new formalisms, tools, techniques, and development methodologies trying, on the one hand, to facilitate the way long-lived successful software systems can be changed in order to cope with demands from users and the increasing complexity and volatility of the contexts in which such systems operate, and, on the other hand, to understand and if possible control the processes by which demand for these changes come about.

The 2010 edition of the workshop will be held in conjunction with the 25th IEEE/ACM International Conference on Automated Software Engineering (ASE'2010), which takes place on 20-24 September 2010 in Antwerp, Belgium. Hence, the special theme of the workshop will be automation in the context of software evolution. Papers should preferably present concepts, techniques or methodologies that are automated or amenable to automation.

Topics of interest for the workshop include, but are not limited to:
- Application areas: distributed, embedded, real-time, ultra large scale, and self-adaptive systems, web services, mobile computing, information systems, systems of systems, etc.
- Paradigms: support and barriers to evolution in aspect-oriented, agile, component-based, and model-driven software development, service-oriented architectures, etc.
- Technical aspects: co-evolution and inconsistency management, impact analysis and change propagation, dynamic reconfiguration and updating; architectures, tools, languages and notations for supporting evolution, etc.
- Managerial aspects: effort and cost estimation, risk analysis, software quality, productivity, process support, training, awareness, etc.
- Empirical studies related to software evolution.
- Industrial experience on successes and failures related to software evolution
- Interdisciplinary approaches: adaptation of evolutionary concepts and measures from other disciplines (biology, geology, etc.) to software evolution
- Theories and models to explain and understand software evolution

Important dates:
- Abstract due: 4 June 2010
- Paper due: 11 June 2010
- Notification: 23 July 2010
- Workshop: 20-21 September 2010

More information:
IWPSE-EVOL'2010:
http://ssel.vub.ac.be/iwpse-evol/

ASE'2010:
http://soft.vub.ac.be/ase2010/

ERCIM WG on Software Evolution:
http://wiki.ercim.eu/wg/SoftwareEvolution/

Please contact:
Tom Mens
Institut d'Informatique, Université de Mons/FNRS/FWO, Belgium
E-mail: iwpse-evol@evolumons.be

sponsored by ERCIM

CSL 2010 - Annual Conference of the European Association for Computer Science Logic

Brno, Czech Republic,
23-27 August 2010

Computer Science Logic (CSL) is the annual conference of the European Association for Computer Science Logic (EACSL). The conference is intended for computer scientists whose research activities involve logic, as well as for logicians working on issues significant for computer science. The 19th EACSL Annual Conference on Computer Science Logic (CSL 2010) and the 35th International Symposium on Mathematical Foundations of Computer Science (MFCS 2010) are federated and organized in parallel at the same place. The federated MFCS & CSL 2010 conference has common plenary sessions and social events for all participants. The technical program and proceedings of MFCS 2010 and CSL 2010 are prepared independently. The MFCS & CSL 2010 conference is accompanied by satellite workshops on more specialized topics.

Topics of interest include (but are not limited to) automated deduction and interactive theorem proving, constructive mathematics and type theory, equational logic and term rewriting, automata and games, modal and temporal logic, model checking, decision procedures, logical aspects of computational complexity, finite model theory, computational proof theory, logic programming and constraints, lambda calculus and combinatory logic, categorical logic and topological semantics, domain theory, database theory, specification, extraction and transformation of programs, logical foundations of programming paradigms, verification and program analysis, linear logic, higher-order logic, nonmonotonic reasoning.

Proceedings will be published in the Advanced Research in Computing and Software Science (ARCoSS) subline of the LNCS series.

There will be a Special Issue of the Journal LMCS (Logical Methods in Computer Science) based on selected papers of CSL 2010.

CSL/MFCS Plenary Speakers
- David Basin (Zurich)
- Herbert Edelsbrunner (Klosterneuburg)
- Erich Graedel (Aachen)
- Joseph Sifakis (Gieres)

CSL Invited Speakers
- Peter O'Hearn (London)
- Jan Krajicek (Prague)
- Andrei Krokhin (Durham)
- Andrey Rybalchenko (Munich)
- Viktor Kuncak (Lausanne)

More information:
http://mfcsl2010.fi.muni.cz/csl

Call for Participation

CEDI 2010 - Third Spanish Conference on Informatics

Valencia, Spain, 7-10 September 2010

Several national symposia and workshops have traditionally been organized on specific topics related to computer science and engineering. CEDI joins all of them into a single conference, and this is the third time that CEDI is organized. The main goal is to gather the Spanish research community on computer science and engineering in order to proudly show our scientific advances, discuss specific problems and gain visibility in Spanish society, emphasizing Spain's role in the new age of the Information Society.

The conference is organized in a 'federated conference' format, by joining several more specific symposia (22 and growing), including Artificial Intelligence, Software Engineering and Databases, Programming Languages, Parallelism and Computer Architecture, Computer Graphics, Concurrency, Soft-computing, Pattern Recognition, Natural Language Processing, Ubiquitous Computing and Human-Computer Interaction. In addition to the scientific activities of each symposium, the CEDI conference will feature two invited keynotes, two round tables and social events and a gala dinner at which several awards will be delivered. Through these awards, the research community on informatics will recognize the efforts of some colleagues in promoting informatics research in Spain.

The expected attendance is around 1,500 scientists. The event is funded by the Spanish Ministry of Education, and sponsored by SpaRCIM. The conference will be organized by the Universidad Politécnica de Valencia. The conference chair will be Prof. Isidro Ramos, the Scientific Program Chair will be Prof. José Duato, and the Local Organization Chair will be Prof. Juan Miguel Martinez, all of whom are from the Universidad Politécnica de Valencia. This university is one of the leading technical universities in Spain, and has a strong focus on innovation and research.

More information:
http://www.congresocedi.es/2010/

Call For Participation

ETM 2010 - 3rd Workshop on Economic Traffic Management

Amsterdam, The Netherlands,
6 September 2010
collocated with 22nd International Teletraffic Congress

The 3rd Workshop on "Economic Traffic Management (ETM)" will take place on September 6, 2010, Amsterdam, The Netherlands. It is the continuation of two successful events that were held at the University of Zurich in the years of 2008 and 2009. The main objective of ETM 2010 (supported by the FP7 STREP SmoothIT) is to give scientists, researchers, and operators the opportunity to present innovative research on ETM, to discuss new related ideas and directions, and to strengthen cooperation in this field of economics-technology interplay. Being collocated with ITC22 – the International Teletraffic Congress - the Workshop on ETM will bring together a new and fast-growing scientific community.

More information:
http://www.csg.uzh.ch/events/etm.

Call For Participation

Workshop on Future Standards for Model-Based User Interfaces

Rome, 13-14 May 2010
Hosted by the CNR-ISTI HIIS Laboratory

Web application developers face increasing difficulties due to wide variations in device capabilities, in the details of the standards they support, the need to support assistive technologies for accessibility, the demand for richer user interfaces, the suites of programming languages and libraries, and the need to contain costs and meet challenging schedules during the development and maintenance of applications.

Research work on model-based design of context-sensitive user interfaces has sought to address the challenge of reducing the costs for developing and maintaining multi-target user interfaces through layered architectures that separate out different concerns.

The goal of this workshop is to discuss the main results of the Model-Based UI XG and to identify opportunities and challenges for new open standards in the area of Model-Based User Interfaces.

Workshop participants will collectively help to identify opportunities and challenges for new open standards in the area, particularly concerning the semantics and syntaxes of task, abstract and concrete user interface models. In addition, workshop participants will have the opportunity to discuss the role of model-based approaches in relation to other standards, for instance, XForms, ANSI/CEA-2018 and MDA, and the relationship to work on standards for Web delivery including HTML5 and browser scripting APIs.

You should consider participating in this workshop if you are in one of the following communities:

- Model-Based (or Model-Driven) UI technology vendors and developers, including open source projects;
- companies that own solutions for the development of multi-target, context-sensitive, adaptive User Interfaces;
- companies seeking to exploit Model-Based Approaches internally;
- software vendors or open source projects that currently offer XML-based languages for the description of user interfaces targeted to the desktop environment;
- government organizations seeking to the standardization of UI development;
- academic researchers with an interest in Model-Based UI Development.

More information:
http://www.w3.org/2010/02/mbui/cfp.html

Call for Submissions and Participation

Modeling Wizards 1st International Master Class on Model-Driven Engineering

Oslo, 30 September - 2 October 2010
A special edition of the MDD4DRES International School co-located with the MODELS Conference

Over the course of the years, computer science has reached a level of maturity that requires an ever greater use of modeling. For over a decade, the MODELS conference has served as the premier venue for the exchange of innovative technical ideas and experiences relating to the use of model-based approaches in the development of complex software systems. During this time, many of its keynote speakers have pointed out the critical role that education plays in expanding the use of models in software development. The Modeling Wizards Master Class, organized jointly with the MODELS conference, offers its participants a set of carefully selected lectures dealing with various aspects of modeling, with a particular focus on domain-specific languages and on the practical application of model-driven tecninques to concrete scenarios.

All of our lecturers are well-known experts, from academia or industry. The academic lecturers will provide an overview of their areas of research, covering the state of the art and outlining emerging trends and challenges. The contributions of the industrial lecturers will focus on results achieved in the practical application of modeling in industrial contexts.

Important Dates:

- Early registration deadline: July 1st, 2010 (subject to limitations).
- Poster submission deadline: August 1st, 2010
- Registration fee is 6400 NOK (about 800 EUR).

Scientific Advisory Board

- Ileana Ober, IRIT – University of Toulouse, France
- Sébastien Gérard, CEA France
- Robert France, Colorado State University, USA
- Oystein.Haugen, SINTEF, NO
- Alexander Pretschner, Technische Universität Kaiserslautern, Germany
- Andy Schuerr, Technische Universität Darmstadt, Germany
- Bran Selic, Canada
- Naoyasu Ubayashi, Kyushu Institute of Technology, Japan
- Jon Whittle, Lancaster University, UK

Organizing Comittee

- Ileana Ober, IRIT – University of Toulouse, France
- Sébastien Gérard, CEA France
- Antonino Sabetta, CNR-ISTI, Italy

More information:
http://modelingwizards.isti.cnr.it/
http://www.mdd4dres.info/
http://www.modelsconference.org/

Call For Participation

ICT 2010

27-29 September 2010

ICT 2010, Europe's biggest research event for information and communication technologies, will be organised by the European Commission DG Information Society and Media on 27-29 September 2010 in Brussels. ERCIM and ERCIM-managed projects will most certainly have representation at this key event through booths and speakers, as in Lyon, France, in November 2008.

More information:
http://ec.europa.eu/information_society/events/ict/2010/

Meet4cleantech – Cooperation in Clean Technologies

Geneva, Switzerland, 1-2 June 2010

The Cleantech Cluster of Western Switzerland will be launched at international level and the European Partnering Event for Clean Technologies will be held in Geneva. Through a series of face2face meetings, representatives of SMEs, industry and research centres from throughout Europe will initiate commercial and technological cooperation and European Research & Development projects in three areas of Clean Technologies: Energy Efficient Buildings, Green Cars and Factory of the Future.

The meet4cleantech event is initiated by Euresearch and organised in cooperation with the Enterprise Europe Network in Switzerland, Osec, the Cleantech Cluster Western Switzerland and OPI to provide a "marketplace" for partners bringing together companies and research centres from throughout Europe. The aim is to facilitate partnerships development, accelerating innovation and initiating research projects. Participants will provide profiles with information on their respective technologies, products and services for inclusion in an online catalogue. This will enable participants to identify potential partners that have complementary know-how before the event and arrange face2face meetings during meet4cleantech. Thus, networking becomes efficient and effective.

The Partnering Event will cover the following themes: "Energy Efficient Buildings", "Green Cars" and "Factory of the Future". Calls for proposals in these themes covering applied research and development projects are scheduled for publication by the European Commission in July 2010. The themes are part of the "Innovation" pillar of the European Economic Recovery Plan. As this pillar is part of the Seventh Framework Programme, Swiss companies and institutions can participate fully and so are also eligible for funding. The European Partnering Event for Clean Technology will focus not only on partners for research projects but also on direct business and technological cooperation between companies and between companies and research institutes.

New cluster and new export platform
The new Cleantech cluster of western Switzerland will be officially launched at international level on 1st June 2010 in Geneva. This cluster was initiated by 7 Swiss Cantons and the directors of their respective economic departments. The aim is to provide greater visibility to clean technology companies both nationally and internationally and to facilitate cooperation. Hence, a number of sessions on exports and innovation will be held on 1st June: The Osec will present it's services for SMEs to access new markets via the new Cleantech export platform. Futhermore, innovative SMEs will be able to get an overview of the initiatives available for supporting innovation at Swiss and international levels.

More information:
http://www.meet4cleantech.eu

HPC Center Pisa Workshops: XtreemOS for Providing HPC Services

Pisa, Italy, 16 June 2010

XtreemOS is a new LINUX-based operating system, improved with Grid and Cloud oriented services to provide the abstraction of a geographically distributed computing platform which is pervasive, secure and easy to use.

XtreemOS is the result of the EU-funded project FP6-033576. It integrates into a single computing platform a wide range of LINUX devices, ranging from ordinary workstations to Single-system-image HPC clusters, and including mobile devices as well. XtreemOS is designed as a scalable and flexible system, fit for industrial use and the ubiquitous needs of Future Internet.

The use of XtreemOS on top of Grid and Cloud computing resources will be the topic of the upcoming workshop "X-HPC, XtreemOS for providing HPC services", hosted at the Department of Computer Science by the new-born HPC center of the University of Pisa.

The HPC center of the University of Pisa experiments with the latest and bleeding edge technology in hardware and software, like GPGPU technologies on ATI and NVIDIA devices, and has been nominated the first ACER HPC competence center in Europe.

The X-HPC workshop will be held in Pisa, on Wednesday 16 June, in the Gerace room at the Computer Science Department of the University. Several industrial partners of the XtreemOS consortium and of the HPC center, like Kerlabs and Ferrari Racing Division, will join the event to provide their unique expertise and viewpoint. The workshop will showcase the latest XtreemOS LINUX distribution, which is currently under testing and evaluation on top of a network of computing resources provided by the Pisa HPC center, by INRIA, the Aladdin/G5K project, CNR and several other partner institutions of the XtreemOS consortium across all Europe. Further details about the event are available on the XtreemOS web site.

More information:
http://www.xtreemos.eu

TPAC 2010

Lyon, France, 1-5 November 2010

A week of combined Technical Plenary / Advisory Committee meetings will be held in Lyon, France at "Cité Centre de Congrès de Lyon". Over 300 of the Web's leading technical experts from about 200 of W3C's Member organizations regularly attend this annual week of W3C meetings. The Technical Plenary meeting itself (on Wednesday 3 November) provides an excellent preview for companies that have an interest in participating in W3C work, while appreciating the strong and unique technical environment that has characterized W3C's success.

Jan Flusser, Tomáš Suk and Barbara Zitová

Moments and Moment Invariants in Pattern Recognition

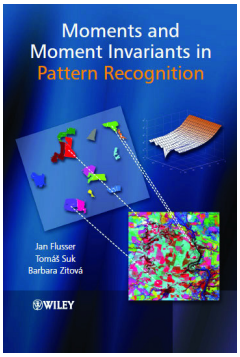

The book presents a unique overview of recent as well as traditional image analysis and pattern recognition methods based on image moments. Invariants to traditional transforms - rigid-body and affine - are studied in-depth from a new point of view. Recent results on invariants to linear filtering of the image and on implicit moment invariants to elastic deformations are clearly presented and well explained. Various classes of orthogonal moments are reviewed and compared and their application to image reconstruction from moments is demonstrated. The authors review efficient numerical algorithms that can be used for moment computation in a discrete domain. Finally, many practical examples of using moment invariants are demonstrated.

Wiley & Sons Ltd., 2009 (312 pp., ISBN 978-0-470-69987-4)

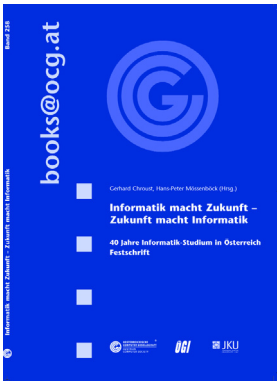

40 Years of Informatics Curricula at Austrian Universities

Fourteen years ago, on August 13, 1969 the Austrian Parliament passed a new law on university studies which for the first time included Informatics curricula. A Festschrift (in German) has now been published which not only summarizes the celebration at the University Linz in April 2009, but additionally includes 30 contribution written by eminent contemporary witnesses, and pioneers together with interviews with several key persons from academia and industry. It also provides a survey of the current status of informatics curricula in Austria.

Order information:
Gerhard Chroust, Hans-Peter Mössenböck (Hrsg.): „Informatik macht Zukunft – Zukunft macht Informatik - 40 Jahre Informatik-Studium in Österreich – Festschrift“ (in German). Schriftenreihe der Österreichischen Computer Gesellschaft, books@ocg, Band 258, November 2009
ISBN 978-3-85403-258-8, price € 25.- from:
Österreichische Computer Gesellschaft
Wollzeile 3-5, 1010 Wien
E-Mail: ocg@ocg.at
Tel: +43 1 5120235 , Fax: +43 1 5120235 9
http://www.ocg.at/publikationen/

The 2009 Tony Kent Strix Award Winner is Carol Ann Peters

The Tony Kent Strix Award, given by the UK eInformation Group of CILIP: the Chartered Institute of Library and Information Professionals, has been awarded to Carol Ann Peters of Istituto di Scienza e Tecnologia dell'Informazione “A. Faedo”.

Carol was nominated for her work on the Cross Language Evaluation Forum (CLEF), the world's leading forum for evaluating cross language searching systems, which Carol initiated and has run for ten years. This evaluation exercise has attracted a multi-disciplinary network of researchers to collaborate on shared tasks and to meet annually to present and discuss results. Thanks to Carol's hard work, CLEF has expanded year over year to become a major international event in information retrieval related research, attracting participation from leading research groups from Europe, the Middle East, Asia and the Americas. Carol's tireless work with CLEF has driven the information retrieval community to consider search in more than just the English language that previously dominated information retrieval research; thus she has made an enormous contribution to the awareness and understanding of information retrieval.

It is difficult to encapsulate here the huge impact of CLEF. Multilingual search is garnering more interest in professional circles, where previously nationally-based work tasks have been extended to have international scope. Carol’s contributions to information retrieval have been many and wide ranging, covering all of the areas for which the Tony Kent Strix Award is offered. The Tony Kent Strix Award is presented each year in memory of Dr Tony Kent, a past Fellow of the Institute of Information Scientists, who died in 1997. Tony Kent made a major contribution to the development of information science both in the UK and internationally. The award is offered in recognition of individuals or groups for an outstanding contribution, practical innovation or achievement in the field of information retrieval.

Scenes from CLEF 2009 Workshop.

Domenico Laforenza appointed ERCIM BoD representatve for CNR

Domenico Laforenza has been appointed as the representative of the Italian National Research Council (CNR) on the ERCIM Board of Directors, following the resignation of Professor Francesco Beltrame. Dr Laforenza is the Director of the "Istituto di Informatica e Telematica" (IIT-CNR) in Pisa. He is also head of the Registry for domain names for Italy (ccTLD.it), and member of the Board of Directors of EURid, the European Registry of Internet Domain Names. In March 2009 Dr. Laforenza was appointed as Italian representative to the European Union "Future Internet Forum of the Member States" by the Italian Ministry of Research, Education and University. From 2006-2009, he was member of the Scientific Advisory Board of the CNR Department for Information and Communication Technologies. and is currently member of the Scientific Council of the Institute de Grilles of the French CNRS

Dr Laforenza has participated as CNR scientific coordinator in numerous European ICT projects (FP6: CoreGrid, NextGrid, GridCoord, XtreemOS, GridComp, Challengers, and FP7: S-CUBE) and was chairman of the Members General Assembly (MGA) of "CoreGrid: the European Research Network on Foundations, Software Infrastructures and Applications for large scale distributed, GRID and Peer-to-Peer". Since June 2006, he has been chairman of the Governing Board of the European project "XtreemOS".

His main research interests are in high performance computing, parallel and distributed systems programming, Grid and Internet computing and he was Head of the High Performance Computing Laboratory at the CNR Institute of Information Science and Technologies (ISTI-CNR) until 30 June 2008.

RESERVOIR Leading the Future of Cloud Computing

The European Commission (EC) recently highlighted, in an Expert Group's Report on the Future of Cloud Computing, the need for coordinated open source deployments between Research and Industry. This is to promote the use of flexible commercial Cloud-based services infrastructure offerings. The RESERVOIR (Resources and Services Virtualization without Barriers) European Framework Programme (FP7) project, now in its third year, has developed open architecture and interfaces, and open source software. These are available as downloadable specifications and tools encompassed under "the RESERVOIR Framework" which presents a blueprint to help businesses build on-demand infrastructure services, at competitive costs, across disparate administrative domains, while assuring quality of service.

The RESERVOIR approach, led by a group of strong industrial and academic partners is focusing on the federation of Clouds at the infrastructure level, to enable on-demand delivery of IT services at competitive costs without requiring a large capital investment in infrastructure. With this approach, end-users can run their applications on virtual machines over the cloud, and are relieved of the need to own and manage physical resources. The software and specifications that are being released by the project offer users an opportunity to build public and private clouds by choosing specific components.

Interoperability is important to users and RESERVOIR is working closely on standard interfaces for the remote management of cloud infrastructures, working with various standards bodies including DMTF and OGF between the Virtual Infrastructure Management and the Cloud Management layers. RESERVOIR is collaborating with several related projects. One of these is OpenNebula, an open-source software application toolkit for building Cloud deployments. OpenNebula is one of the technologies being enhanced in the RESERVOIR project as a service solution for Cloud management.

More information:
http://www.reservoir-fp7.eu/

Jeffrey Jaffe Named W3C CEO

W3C named Dr. Jeffrey Jaffe its new Chief Executive Officer. "Web technologies continue to be the vehicle for every industry to incorporate the rapid pace of change into their way of doing business," said Dr. Jaffe. "I'm excited to join W3C at this time of increased innovation, since W3C is the place where the industry comes together to set standards for the Web in an open and collaborative fashion." As W3C CEO, Dr. Jaffe will work with Director Tim Berners-Lee, staff, Membership, and the public to evolve and communicate W3C's organizational vision. The CEO is responsible for W3C's global operations, for maintaining the interests of all of the W3C's stakeholders, and for sustaining a culture of cooperation and transparency, so that W3C continues to be the leading forum for the technical development and stewardship of the Web.

ERCIM
European Research Consortium for Informatics and Mathematics

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

SpaRCIM
Spanish Research Consortium for Informatics and Mathematics
Spanish Research Consortium for Informatics and Mathematics,
D3301, Facultad de Informática, Universidad Politécnica de Madrid,
Campus de Montegancedo s/n,
28660 Boadilla del Monte, Madrid, Spain,
http://www.sparcim.es/

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

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

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

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

Order Form

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

Send, fax or email this form to:

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

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

I wish to subscribe to the

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

Name:

Organisation/Company:

Address:

Postal Code:

City:

Country

E-mail:

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