

ERCIM NEWS

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

Number 70, July 2007

Special: Service-Oriented Computing



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

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

Editorial Board:

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

Local Editors:

Austria: Erwin Schoitsch, (erwin.schoitsch@arcs.ac.at)
Belgium: Benoît Michel (michel@tele.ucl.ac.be)
Czech Republic: Michal Haindl (haindl@utia.cas.cz)
Finland: Pia-Maria Linden-Linna (pia-maria.linden-linna@vtt.fi)
France: Bernard Hidoine (bernard.hidoine@inria.fr)
Germany: Michael Krapp (michael.krapp@scai.fraunhofer.de)
Greece: Eleni Orphanoudakis (eleni@ics.forth.gr)
Hungary: Erzsébet Csuhaaj-Varjú (csuhaj@sztaki.hu)
Ireland: Ray Walsh (ray@computing.dcu.ie)
Italy: Carol Peters (carol.peters@isti.cnr.it)
Luxembourg: Patrik Hitzelberger (hitzelbe@lippmann.lu)
Norway: Truls Gjestland (truls.gjestland@ime.ntnu.no)
Poland: Hung Son Nguyen (son@mimuw.edu.pl)
Spain: Salvador Lucas (slucas@dsic.upv.es)
Sweden: Kersti Hedman (kersti@sics.se)
Switzerland: Harry Rudin (hrudin@smile.ch)
The Netherlands: Annette Kik (Annette.Kik@cwi.nl)
United Kingdom: Martin Prime (M.J.Prime@rl.ac.uk)
W3C: Marie-Claire Fogue (mcf@w3.org)

Contributions

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

Copyright Notice

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

Advertising

*For current advertising rates and conditions, see
<http://ercim-news.ercim.org/> or contact office@ercim.org*

ERCIM News online edition

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

Subscription

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

Next issue:

*October 2007
Special theme: Technology-Enhanced Learning*

The evolution towards a service-based economy represents today more than 70% of the GDP in developed countries. This evolution is expected to continue and to influence our lifestyle in the coming decades but the full potential of the service-based economy can only be grasped if it builds on new technologies for the design, development and operation of new services and service platforms.

This trend towards a service-based economy is a challenge to everybody and a new opportunity for Europe to position itself at the centre of innovation in this global race. European research and industry should play a key role in the establishment of the future IT service platforms, notably by influencing global standards, and contribute to the development of the capabilities required for Europe to design and deploy innovative services in order to remain competitive in our new globalised and knowledge-based world.

Over the past years, Service-Oriented Architectures have become a paradigm among enterprises for enabling more efficient and flexible business processes and addressing some of the technological challenges posed by the service-based economy. Using loosely coupled services, they allow discovering, orchestrating and composing capabilities that are needed by end-users, business processes or other services. Service-Oriented Architectures can help businesses to react more quickly and cost-effectively to changing market conditions.

In parallel with these developments, research on Grid Technologies has evolved from its initial focus on complex eScience applications towards general-purpose service infrastructures that can also be used by business and industry. The emphasis is now on the dynamic provision of resources in an easy and transparent manner to allow translating business or user requirements to infrastructure capabilities. The focus is also on the convergence with the service paradigm so that the infrastructure can be managed according to more flexible and dynamic business practices.

This convergence has extended the scope of research on Service-Oriented Architectures to address the implementation of these services in the available Grid infrastructure. The aim is that it becomes a utility allowing the dynamic allocation and assembly of resources as required by the applications, thus supporting business agility and flexibility at the infrastructure level.

A number of research and technical challenges still need to be addressed before this vision becomes reality. To mention just a few, security and privacy are essential factors for take up, but also performance and quality of service, easy manageability and low deployment costs, particularly in the case of large scale systems. Key challenges are also of a business nature: precise accounting across all layers, enforcement and tracking of policies down to the infrastructure level, business model identification, or supporting business agility and more efficient business processes.

The IST research programme of the European Commission has strongly contributed to these developments. During the 6th Framework Programme, the activities related to software and services have catalysed European research on service architectures and related technologies, while the activities on Grid technologies have prepared the transition towards service infrastructures.

In parallel to these activities, the Networked European Services and Software Initiative (NESSI) has brought together key European industrial players in the area to address the major changes that are driving the IT services marketplace. NESSI is developing a strategic research agenda for the software and services area, is addressing the technological, business and policy challenges, and is securing industrial leadership in these developments and their take up.

This provides a good starting point for Europe to continue at the forefront of these developments while being ready to grasp their business and commercial benefits. In the 7th Framework Programme that has just started, Service-Oriented Architectures is one of the pillars of our R&D activities and will provide new opportunities to make the Software and Services vision a reality.

See also: http://cordis.europa.eu/fp7/ict/ssai/home_en.html



*Jesús Villasante,
European Commission,
DG Information Society and Media,
Head of Software & Service
Architectures and Infrastructures Unit.*

2 Editorial Information**3 Keynote**

by Jesús Villasante, European Commission, Head of Unit, Software & Service Architectures and Infrastructures

JOINT ERCIM ACTIONS**6 MUSCLE at CeBIT****6 First ERCIM Workshop on eMobility**

by Torsten Braun

7 ERCIM Working Group 'Dependable Embedded Systems' - Forthcoming Activities and Workshops

by Bruno Le Dantec

8 ERCIM IM2IM Working Group Annual Meeting

by Marc Thiriet

NEWS FROM W3C**9 Calls for Participation in W3C Workshops****9 W3C Workshop Reports****9 Semantic Web Case Studies and Use Cases Published****9 Latest W3C Recommendations****THE EUROPEAN SCENE****10 Photonics is Given a Firm Place in the 7th Framework Programme****11 Second EU-US Workshop on Secure, Dependable and Trusted ICT Infrastructures**

by Jim Clarke and Thomas Skordas

12 Science Foundation Ireland funded Online Dublin Computer Science Summer School**13 EuroSpaceward – Preserving and Improving Life on Earth by Going into Space**

by Markus Klettner, Benoît Michel and Bradley Edwards

SPECIAL THEME**Introduction to the Special Theme****14 Service-Oriented Computing**

by Jana Koehler and Gustavo Alonso

Grid Overview**16 The UK National Grid Service**

by Katie Weeks

17 Mobile Service Management in Service-Oriented Grids

by Tom Kirkham, Fredrik.Solsvik and Robert Piotter

18 A4C Support for Commercialization of Next-Generation Grid Services

by Hasan, Peter Racz, Cristian Morariu, David Hausheer and Burkhard Stiller

20 Collective Completion: How do YOU Complete your Unknown Colleagues?

by Magnus Boman, Kristofer Franzén and Fredrik Espinoza

21 Workflow Management Systems for Grid Computing

by Carmen Bratosin and Wil van der Aalst

Quality of Service and Formal Approaches**22 Towards a Commercial IT Service Delivery**

by Christophe Ponsard, Gautier Dallons, Stéphane Mouton and Philippe Massonet

24 Reliable and Inexpensive QoS Monitoring in Service Markets

by Radu Jurca, Walter Binder, Boi Faltings

25 WS-DIAMOND: An Approach to Web Services – Monitoring and Diagnosis

by Luca Console and Mariagrazia Fugini

27 Formal Modelling and Verification in Service-Oriented Computing

by Maurice ter Beek, Stefania Gnesi, Fabio Martinelli, Franco Mazzanti and Marinella Petrocchi

28 Challenges in a Service-Oriented World

by Wolfgang Reisig, Karsten Wolf, Jan Bretschneider, Kathrin Kaschner, Niels Lohmann, Peter Massuthe, and Christian Stahl

Applications**30 RCE and SESIS - Service-Oriented Integration Environment for Collaborative Engineering**

by Ottmar Krämer-Fuhrmann

31 The ITIA-CNR Customized Shoe Production Service

by Andrea Chiodi, Fabrizio Silva and Andrea Ballarino

32 The ToolBus: A Service-Oriented Architecture for Language-Processing Tools

by Paul Klint

- 34 Open Grid Services for Improving Medical Knowledge Discovery**
by Manolis Tsiknakis
- 35 Application of SensorWebs for Intelligent Satellite Tasking**
by Adrian Grenham
- 37 Open-Source SOA for Small and Medium Organisations**
by Pascal Bauler and Nicolas Biri
- [Semantic Web, SOA, Business Processes](#)
- 38 FUSION: Business Process Fusion Based on Semantically Enabled Service-Oriented Business Applications**
by Spiros Alexakis and Alexa Schumacher
- 40 VantagePoint: Seeing is Believing**
by Jarmo Kalaoja, Julia Kantorovitch, Toni Piirainen and Ilkka Niskanen
- 41 Diapason: An Engineering Environment for Designing, Implementing and Evolving Service Orchestrations**
by Frédéric Pourraz and Hervé Verjus
- 43 SUPER – Raising Business Process Management Back to the Business Level**
by Matthias Born, Christian Drumm, Ivan Markovic and Ingo Weber
- 44 Life-Cycle Support of Semantic Web Services**
by László Kovács, András Micsik and Tomás Pariente
- 46 Aligning Virtual Organizations with SOA**
by Mikko Salonen and Jyrki Haajanen
- 47 A Pervasive, Service-Oriented Architecture for Supporting Teamwork**
by Schahram Dustdar and Hong-Linh Truong
- 48 Engineering Services**
by Mike P. Papazoglou
- 49 Business Process-Driven Service Architecture**
by Veronica Gacitua and Claus Pahl
- 51 A Service-Oriented Middleware for B3G Networking**
by Valerie Issarny, Mauro Caporuscio, Pierre-Guillaume Raverdy
- 53 Exploiting Organizational Models for Semantic Service Description, Matchmaking and Composition in Service-Oriented Multi-Agent Systems**
by Alberto Fernández and Sascha Ossowski
- 54 Building Service-Oriented Web Applications for Business Processes**
by Victoria Torres, Vicente Pelechano and Pau Giner

R&D AND TECHNOLOGY TRANSFER

- 56 Coordination of Emergency Communication in Safe Hands**
by Annette Kik
- 57 Super Progress in Super Computing**
by Harry Rudin
- 58 Grid Technology Makes Weather Forecasts without Boundaries a Reality**
by Baudouin Raoult, Clemens-August Thole and Ute Gärtel-Zafiris
- 60 Meeting the Multicore Challenge: Dynamic Dependence Analysis for Parallelization**
by Karl-Filip Faxén
- 61 Multimodal Multi-Level Fusion using Contextual Information**
by Olga Vybornova, Monica Gemo and Benoit Macq
- 62 The SAPIR Project: Executing A/V Complex Queries in Peer-to-Peer Systems**
by Claudio Gennaro, Raffaele Perego and Fausto Rabitti
- 63 Getting your Data Back by Giving it Away**
by Jacco van Ossenbruggen
- 64 OPTIHPER: A Computer-Based Decision Support System for Employee Timetabling Problems**
by Antonio Lova, Pilar Tormos and Federico Barber

EVENTS

- 66 First EchoGRID-EUChinaGRID International Conference**
- 67 Announcements**
- 71 In Brief**

MUSCLE at CeBIT

The MUSCLE Network of Excellence (Multimedia Understanding through Semantics, Computation and Learning), coordinated by INRIA and ERCIM, presented its activities and showcases on a booth at CeBIT, the largest IT fair worldwide, from 15-21 March 2007 in Hannover, Germany. As part of its dissemination activities, various partners of the network presented their multimedia expertise through interactive demonstrations.

Live demos allowed visitors to try out and experiment with 'Maps of Music' exploring musical landscapes and listening to selected music using the PlaySOM application and the 'PocketSOMPlayer' developed at Technical University Vienna. By applying self-organising maps (SOM) - a sub-



MUSCLE booth at CeBit.

type of artificial neural networks - to automatically analyzed audio content (frequency spectra), researchers have created an innovative way to organize music by sound similarity. After categorization, the content of an audio collection is displayed on the screen of a laptop or mobile device as a topographic map where colour represents the density of a particular genre in the collection. The user can then select a playlist reflecting his or her mood by drawing a path through this audio landscape.

A research group at the EC3 - E-Commerce Competence Center, Vienna, Austria, demonstrated 'The Media Square', a SOM-based 3D Virtual World allowing users from all over the world listening to music together as well as attending meetings in a virtual world.

INRIA presented a video-clip copy detection software that identifies copied parts in a video even if it was remastered. This is a useful tool to control the copyrights of digital contents on the web and on TV, and also a powerful tool to analyze and explore these data. A crucial difficulty is the fundamental difference between a copy and the notion of similar image encountered in Content-Based Retrieval: a copy is not an identical or a near replicated video sequence but rather a transformed video sequence. These photometric or geometric transformations (gamma and contrast transformations,

overlay, shift, etc) can greatly modify the signal, and therefore a copy can in fact be visually less similar than other kinds of videos that might be considered similar.

MUSCLE aims at establishing and fostering closer collaboration between research groups in multimedia data mining and machine learning. The goal is to explore the full potential of statistical learning and cross-modal interaction for the (semi)-automatic generation of robust meta-data with high semantic value for multimedia documents.

Link:

<http://www.muscle-noe.org/>

Please contact:

Patricia Ho-Hune, MUSCLE project manager

ERCIM office

E-mail: patricia.ho-hune@ercim.org

First ERCIM Workshop on eMobility

by Torsten Braun

The first ERCIM Workshop on eMobility has been held at University of Coimbra, Portugal on 21 May 2007 prior to the International Conference on Wired/Wireless Communications. More than 20 participants attended this workshop for which twelve presentations have been selected by the international program committee out of 21 submissions.

After the official opening by the chair of the ERCIM eMobility Working Group Torsten Braun, Prof. Luis M. Correia from Technical University of Lisbon gave a keynote talk on "eMobility: present and future challenges". Prof. Correia is a steering committee member of the eMobility Technology Platform (ETP) and chairs the working group on applications. In his very interesting talk, he presented the activities, vision and goals of the ETP. This triggered intensive discussions about the possible future cooperation among ETP and eMobility WG.

Three technical sessions complemented the programme. The first session chaired by Giovanni Giambene (University of Siena) on "Traffic Engineering and Mobility Management" discussed the issues of adaptivity and mobility in future mobile networks. It has been discussed how Orthogonal Frequency-Division Multiplexing (OFDM) modulation in 4G networks can be made adaptive to optimize performance and minimize bit errors. Another presentation showed an adaptive layered video encoding mechanism over wireless networks using a fuzzy-logic based control algorithm. A further important issue are the effects of mobility on network performance. The number of handovers can be important for system performance parameters such as call blocking probability. The last talk of this session identified severe problems of handover latencies in Mobile IP networks by analysis of protocols in real test-beds. Significant research efforts seem to be required in the future in order to solve these problems.

Another hot research topic has been discussed in the first afternoon session on "Wireless (Sensor) Networks" chaired by Geert Heijenk (University of Twente). Energy-efficiency is a key problem in sensor networks and approaches for communication protocol design mostly rely on cross-layer concepts, where various protocol layers such as routing and medium access control are integrated with each other. Positions of tracked objects and events should be determined as accurate as possible. Several positioning technologies exist. One talk proposed to combine these in order to increase accuracy. Another important issue is security in wireless sensor networks. A variety of security mechanisms exist, but many of them are too complex to be implemented on available sensor node resources. A compiler has been proposed to select and configure appropriate security mechanisms and protocols dependent on application requirements and sensor node constraints.

The final session discussed various issues in "Pervasive Computing and Mobile Applications" chaired by Saverio Mascolo (Poltecnico di Bari). Many context-sensitive applications and services are currently being developed. Context information, eg location or sensor information need to be exchanged. This can be done in different ways, eg by intelligent flooding, shared tuple spaces, or mobile agents. The fourth talk addressed the issue of authentication in a mobile service environment using available technologies from GSM.



Keynot speaker Luis M. Correia together with Edmundo Monteiro, Universtiy Coimbra, the local organizer of the workshop.

The first workshop of the ERCIM WG on eMobility has been a big success in particular when considering the quality of the presentations. This could only be achieved due to the excellent work of authors and reviewers. The hosts at University of Coimbra not only provided excellent meeting facilities, but also organized to published printed workshop proceedings (ISBN: 978-972-95988-9-0).

Link:

<http://emobility.unibe.ch/>

Please contact:

Torsten Braun
 ERCIM eMobility Working Group coordinator
 Berne University / SARIT, Switzerland
 E-mail: braun@iam.unibe.ch

ERCIM Working Group 'Dependable Embedded Systems': Forthcoming Activities and Workshops

This year, the ERCIM Working Group 'Dependable Embedded Systems' (DES-WG) will (co)organize several workshops together with the European Integrated Project DECOS (Dependable Embedded Components and Systems).

DES-WG is a partner of the DECOS project, and is charged with the task of disseminating the project results. In return, DECOS sponsors the DES-WG workshops. DES-WG has a presence at the following events:

- Open Forum on Time-Triggered Technologies; Toulouse, 24-25 April 2007
- ARTEMIS Annual Conference; Berlin, 4 June 2007
- ARTEMIS Standardization Working Group Meeting; Brussels, 25 May 2007
- INDIN07 - Industrial Informatics; Vienna, 23-26 July 2007 (with a special session on dependable embedded systems – see <http://www.indin2007.org>)
- Euromicro 2007; Lübeck, 27-31 August 2007 (with an ERCIM/DECOS workshop (<http://www.euromicro.org>))
- SAFECOMP 2007; Nuremberg, 18-21 September 2007 (with a DECOS/ERCIM workshop – see call for papers on page 69).



ERCIM/DECOS booth

The papers of INDIN07 will be published as IEEE proceedings, while the other workshop papers will also be reviewed and will be published as ERCIM proceedings. The workshops cover major issues of embedded systems technology, and both posters and oral presentations are given. The areas covered include basic technology, overviews, research challenges, application areas, and education and training. They are excellent opportunities to advertise both ERCIM and the DECOS project.

Links:

<http://its.arcs.ac.at/ercim/>
<http://www.decoss.at/>

Please contact:

Erwin Schoitsch, ERCIM DES-WG coordinator
 Austrian Research Centers - ARC / AARIT
 E-mail: erwin.schoitsch@arcs.ac.at

ERCIM IM2IM Working Group Annual Meeting

by Marc Thiriet

The ERCIM Working Group 'IT and Mathematics applied to Interventional Medicine' (IM2IM) held its annual meeting in Lausanne, Switzerland on 14-15 May. The meeting was held during the ERCIM spring meetings and the SARIT'07 symposium on the topic 'Computing for Health' which focused on the Virtual Physiological Human.

The working group (WG) meeting had two sessions: the first on Monday 14 May with presentations from the WG members, and the second dedicated to discussions on possible proposals under the 7th Framework Programme.

The first talk of the presentation session was given by Muriel Boulakia from the INRIA REO (numerical simulation of biological flows) team (in collaboration with M Fernandez, JF Gerbeau and N Zemzemi), who presented results from the CardioSense3D project. This involves modelling of heart activity, including electromechanical coupling, myocardium perfusion and tissue remodelling, as well as estimation of patient-specific parameters from observations, simulations of pathologies and modelling of the effects of therapy. A particular emphasis is given to heart electrophysiology and numerical simulations of ECG signals from the usual twelve leads.

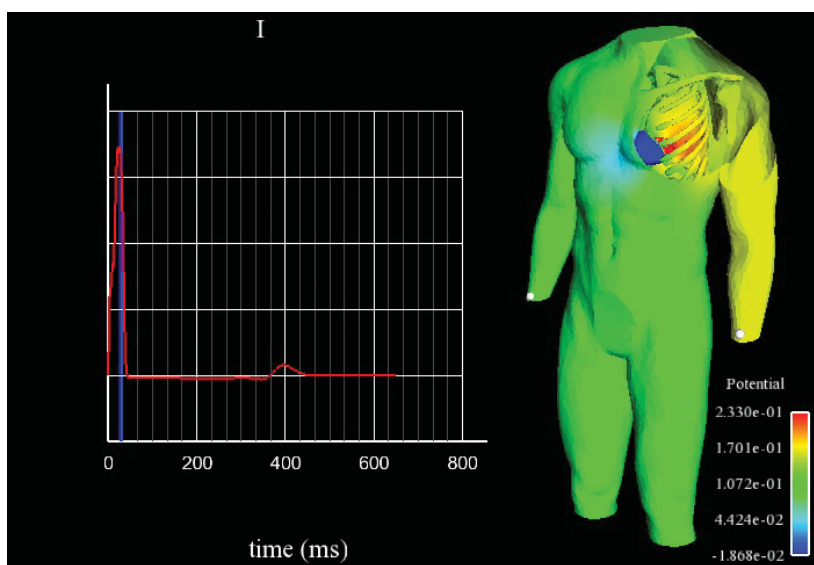
Carlo D'Angelo, of the Institute of Analysis and Scientific Computing from EPFL, presented the second talk, and spoke about 1D-3D modelling of tissue perfusion. Since tissue is perfused by vessels of varying size, perfusion modelling is based on multiscale approaches. Blood flow and mass transport can be described by 1D models in complex networks, but in some cases, 3D models may be necessary. The variational formulation is used for the 1D problem, allowing boundary conditions to be used only at the ends of the explored vasculature. The capillary matrix is modelled by a 3D homogeneous porous medium. The coupling has been developed between the 1D model of the arterial tree, and the 3D homogeneous porous medium that represents the capillary compartment in which the arterial tree is embedded. The pressure distribution, as well as molecule transport through the blood vessels belonging to both the 1D model and capillary bed of the 3D domain, can be computed within the tissue, as can its evolution during the cardiac cycle. Two examples were provided: the retina and the head tissues near the Willis cycle.

Simone Deparis, from the Department of Applied Mathematics at EPFL, in association with Alessandro Veneziani from Politecnico di Milano, spoke on multiscale (0D-1D-3D) modelling of blood flow. The boundary conditions are determined by imposing either the pressure or the flow con-

ditions, with the latter being achieved by a Lagrange multiplier technique. An aneurysm of the internal carotid artery was used to illustrate this work. The pressure distribution in the exit cross-section was not significantly different between the test with stress-free conditions, and the reduced model downstream from the 3D domain coupled to it. However, good estimates of parameters of the lumped parameter model are still lacking.

Erik Burman, from the Institute of Analysis and Scientific Computing at EPFL, in collaboration with M Fernandez, INRIA REO team) gave a talk on stabilized explicit coupling for fluid-structure interaction using Nitsche matching conditions for the variational formulation. Bilinear forms are introduced both for the fluid and the solid dynamics, and a finite element formulation is developed using the Nitsche matching condition. The incompressibility at the interface is relaxed to allow for the mismatch introduction in the coupling. Stabilization for explicit methods can then be obtained.

Finally, Luc Soler from IRCAD spoke on computed assisted surgery starting from fully automatic 3D reconstruction. Augmented reality is then developed using the same tech-



Electrical activity of the heart coupled to the torso: first standard lead of an electrocardiogram. A simulation realised by the INRIA REO team.

nique. Current simulators mostly use artificial organs. Moreover, deformation of organs must be more accurately modelled. Preoperative augmented reality uses the preoperative registration, but does not take into account breathing and associated organ motions. Prediction according to the skin position is necessary, and the aim is to develop more realistic simulations. Besides, ultrasound elastography (fiberscan) allows the distribution of the elastic modulus to be measured. Liver fibrosis must be detected using such a technique.

Please contact:

Marc Thiriet, IM2IM WG coordinator
INRIA, France
E-mail: marc.thiriet@inria.fr

Calls for Participation in W3C Workshops

The goal of a W3C workshop is to convene experts and other interested parties for an exchange of ideas about a technology or policy. These are high profile events with topics related to innovative and visionary Web usage and development.

Next Steps for XML Signature and XML Encryption Mountain View, California, USA, 25-26 September 2007

The XML Signature and XML Encryption specifications have seen broad deployment, and form the basis for a number of security related specifications in the Web and Web Services worlds. This workshop will serve as a means to discuss experience and issues with the XML Security suite of specifications and to identify gaps and emerging issues that might feed future work. Position Papers are due 14 August 2007.

<http://www.w3.org/2007/xmlsec/ws/cfp>

W3C/OpenAjax Alliance Workshop on Mobile Ajax San Francisco Bay Area, USA, 28 September 2007

Position papers are due 15 August for the Workshop on Mobile Ajax co-sponsored by W3C and the OpenAjax Alliance. Attendees will explore use cases for mobile Ajax to help shape its use in mobile Web browsers. Topics might include user experience, application development, support in today's devices and browsers, and whether needs exist for standardization and best practices.

<http://www.w3.org/2007/06/mobile-ajax/>

W3C Workshop on RDF Access to Relational Databases Boston, MA, USA, 25-26 October 2007

Systems for exporting relational data to RDF have existed since the beginning of the Semantic Web. Recently, the semantic Web developers have focused on SPARQL query-rewriters and interpreters to access relational data directly. Both of these approaches share an expression of relational data in RDF. This workshop will draw members of these semantic Web and relational database communities together to examine commonalities, distinctions and next steps for expressing relational data in RDF.

<http://www.w3.org/2007/03/RdfRDB/cfp>

W3C Workshop Reports

A typical W3C workshop's output is the publication of a report which indicates what are the recommended suggestions for the development of future Web technologies. Organizations interested in these published conclusions should get involved in further work by participating in W3C working groups.

"Web of Services for Enterprise Computing" Workshop

The participants recommended, among other things, that W3C should ensure the stability of the Core platform of Web Services by maintaining its Web Services specifications through a Web Services Core Working Group; facilitate the

common accessibility of resources by both Web and Web Service mechanisms; help users/vertical industries to specify needs, use cases, best practices guidelines, and architectural patterns for service-oriented architectures.

http://www.w3.org/2007/04/wsec_report

HTML Mail Workshop

The participants emphasized that the state of HTML implementation in mail user agents lags far behind that of Web browsers, even if the same technology is used for both. HTML email is also subject to additional restrictions caused by security, anti-spam, and anti-phishing initiatives which can negatively impact legitimate HTML email uses.

<http://www.w3.org/2007/05/html-mail/>

Semantic Web Case Studies and Use Cases Published

The Semantic Web Education and Outreach (SWEO) Interest Group is pleased to announce the first set of Case Studies and Use Cases giving some examples of how the Semantic Web of machine readable data is used today. Applications are presented in areas ranging from automotive to health care, and from B2B systems to geographical information systems. Case studies include descriptions of systems that have been deployed within an organization, and are now being used within a production environment. Use cases include examples where an organization has built a prototype system, but it is not currently being used by business functions. As the list is continuously updated, an RSS feed is available to track new additions.

<http://www.w3.org/2001/sw/sweo/public/UseCases/>

Latest W3C Recommendations

- Web Services Description Language (WSDL) Version 2.0 Part 1: Core Language
26 June 2007, Jean-Jacques Moreau, Arthur Ryman, Roberto Chinnici, Sanjiva Weerawarana
- Web Services Description Language (WSDL) Version 2.0 Part 2: Adjuncts
26 June 2007, Sanjiva Weerawarana, David Orchard, Hugo Haas, Roberto Chinnici, Amelia A. Lewis, Jean-Jacques Moreau
- Web Services Description Language (WSDL) Version 2.0 Part 0: Primer
26 June 2007, Canyon Kevin Liu, David Booth
- Voice Extensible Markup Language (VoiceXML) 2.1
19 June 2007, Daniel C. Burnett, Emily Candell, Jerry Carter, RJ Auburn, Scott McGlashan, Paolo Baggia, Alex Lee, Brad Porter, Matt Oshry, Ken Rehor, Michael Bodell, David Burke
- Semantic Interpretation for Speech Recognition (SISR) Version 1.0
5 April 2007, Luc Van Tichelen, David Burke
- Internationalization Tag Set (ITS) Version 1.0
3 April 2007, Christian Lieske, Felix Sasaki

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

Photonics - A Highlight in the 7th Framework Programme

Photonics is one of the most promising and exciting fields currently being developed. Photonics is to light what electronics is to electricity. Or as the eminent French Scientist Pierre Aigrain described it: "Photonics is the science of the harnessing of light. Photonics encompasses the generation of light, the detection of light, the management of light through guidance, manipulation, and amplification, and, most importantly, its utilisation for the benefit of mankind."

The subject was born with the invention of the laser in 1960, which together with the development of optical fibres for communications has led to high capacity telecommunications and is the actual backbone of the internet. Photonics holds a huge potential – not only for new and even better forms of communications and entertainment but also in



Dr. Rosalie Zobel, Director, DG Information Society and Media.

many other applications, including manufacturing, medicine, displays, and a whole range of sensors for chemicals, biological materials and in the environment. Ultimately, photonics even promises to completely replace microelectronics as the technology that computers use to 'think', leading to a huge increase in performance.

Today, some 200,000 people are employed directly in the photonics industry in Europe, and two million other jobs depend on it. The global market for products enabled by photonics is already €150bn per annum and increasing. In Germany alone, there are some 1,000 optics and photonics SMEs employing 36,000 people – a figure expected to grow by more than 40% by the year 2010. But a crucial point is that apart from its economics, photonics is also a vital strategic technology, in which Europe must maintain its expertise and know-how so as to avoid being left behind in a very quickly changing market. It's not enough to be in the photonics race – you've got to be the best.

To better address these challenges, and particularly to help develop a common European approach to photonics research

and development, Europe's photonics community founded the Photonics21 technology platform in December 2005.

With some 700 members in 35 countries (and counting), this industry-led initiative has mobilised the European photonics research actors into a real community, addressing issues of common interest and developing a joint vision. Its first Strategic Research Agenda, delivered in April 2006, was a major input to the development and review of future EU funded research on photonics.

Recognising its importance and huge potential, photonics has been given a prominent position in the Information and Communication Technologies (ICT) part of the EU's 7th Framework Programme for research (FP7). Some €90m have been allocated to fund the basic photonics technologies during 2007-2008 alone, and it is expected that this will increase over the lifetime of the programme. This is in addition to the significant research programmes going on in many European countries such as Germany, which is allocating some €100m to research on Organic Light Emitting Diodes (OLED), or the UK's Technology Programme, which allocated over £55m to photonics related R&D during 2004-2006.

As well as increasing the level of funding, there was a clear need to establish a core team within the Commission dedicated to photonics in FP7. This was not only because of the recognised growing technological and economic importance of the domain and the increased amount of resources, both human and in terms of funding devoted to it, but also in response to the need for a single contact point for the photonics industry and research community. The new Photonics unit will also provide a link to related areas of the Commission. This includes research on applications using photonics such as in the automotive, biomedical, communications, security and displays areas of the ICT programme but also elsewhere, for example, in the Nanosciences, Nanotechnologies, Materials and new Production Technologies programme, and photonics-related research for SMEs.

The person who the Commission has chosen to head the Photonics unit is Thierry Van der Pyl. Thierry has an extensive experience in the European Commission. Previously, he has headed the units for "High Performance Computing", "Microelectronics", "Trust and Confidence" and more recently the "Future and Emerging Technologies" unit.

"We believe we have the components for success: a first-rate existing knowledge base in the research community; a world-class industry which is getting organised; increased national and European research funding; and outstanding cooperation with industry through the Photonics21 technology platform. With these, we can help to ensure a strong, successful and globally competitive European photonics industry, based on excellent and responsive research, for the benefit and wellbeing of everyone. I am convinced that Photonics will be the technology of the 21st century, and this is an area where the efforts of the Commission can make a real difference" Thierry says.

Link:
http://cordis.europa.eu/fp7/ict/photonics/home_en.html

Second EU-US Workshop on Secure, Dependable and Trusted ICT Infrastructures

by Jim Clarke and Thomas Skordas

The 2nd EU-US research workshop on "Cyber Trust: System Dependability and Security" was held in Illinois, USA, in April 2007. This article presents the themes discussed and the main workshop conclusions.

The second EU-US research workshop on "Cyber Trust: System Dependability and Security" was held in Illinois, USA 26-27 April 2007. It was attended by 40 delegates from the EU and the US, along with a few representatives from Canada, Australia and Japan. The event was organised and hosted by the University of Illinois in close co-operation with the US National Science Foundation (NSF), the US Department of Homeland Security (DHS) and the "Security" unit of the European Commission's Directorate General Information Society and Media. The workshop aimed to ensure progressive continuity of the consensus building achieved at the first workshop on the same subject held in November 2006 in Dublin, Ireland (see related article in ERCIM News No 69 of April 2007). The guiding principle was to identify and develop further those research areas that require and will benefit from international collaboration, while examining the structures and mechanisms that could potentially enable and fund the proposed work.

The workshop was structured around two main technical themes: (1) "Architectures, Protocols and Environments for Trust, Security and Dependability (TSD) of future Polymorphic Networked ICT Systems"; and (2) "TSD Attributes and Mechanisms for future Distributed Services and Content, future Overlay Networks and Applications". A transversal topic was also introduced across these two technical themes covering "Test beds, data sets, and models for quantification, evaluation and validation".

At the workshop, there was a broad agreement that the current EU-US cooperation was necessary and should be further elaborated. The workshop also permitted to draft a joint list of research topics relevant for international cooperation and stimulate further discussion on how this could be implemented. The main workshop conclusions are summarised below:

1. Architectures, Protocols and Environments for TSD of future Polymorphic Networked ICT Systems

The session focussed particularly on the security and dependability of future large polymorphic, networked, multi-formed and interdependent ICT systems. Examples of such systems include: the future internet; the internet of "things"; GPS/Galileo; future wireless and mobile systems, RFID and sensor networks (Post IP, Post 3G); mixed mode environments consisting of diverse computing, communication and storage capacities; service-centric, adaptive, heterogeneous and scale-free ambient environments; etc. It was stressed that present security mecha-

nisms based on boundaries and firewall protection mechanisms do not scale up to address the security and dependability of such future complex systems. There is, thus, a need for joint research in developing autonomic, evolvable and adaptive security mechanisms and policies and new cognitive techniques and semantic models for managing the complexity and interdependencies of such ambient systems and their interaction with users. Security and cryptographic mechanisms and protocols must also be scaled down in order to be inserted in small devices having scarce resources, while at the same time ensuring privacy protection.

2. TSD Attributes and Mechanisms for future Distributed Services and Content, future Overlay Networks and Applications

The session addressed the security of applications and services and the protection of data and software systems; the security in specific service application areas requiring collaboration at global scale like financial, health and critical information infrastructures; and, the security in dynamic service coalitions, including the provision of dynamic service virtualisation in polymorphic environments and the trustworthy delivery of atomic or composite services to end users through a contract that provides them with some security guarantees. International collaboration is needed because such service coalitions have elements of dynamism and trust and require the establishment of contracts over unknown and global paths. Such services need to be trustworthy and accountable while they traverse networks spanning over different countries and administrative domains that obey different regulations and policies. These policies have to be interoperable and dialogue and negotiate in real time.

3. Test-beds, Data Sets, and Models for Quantification (Metrics), Evaluation (Techniques) and Validation (Processes)

The following topics were debated in both the above sessions:

Consumer issues: They include the need to build mechanisms that guarantee privacy, traceability, anonymisation and use of pseudonyms of the legitimate users, while at the same time, they permit to locate, track and trace malevolent users at individual, group or organization levels. A new security paradigm needs to be defined that strikes a reasonable balance between a 'Big Brother Society' and a confidence building, privacy and ethics protecting society. Establishing society's confidence in the new digital world would require creating a palpable security environment that enables citizens to control the type and level of protection associated with the digital goods and services they have access to. It would also require developing trust, security and dependability technologies that are unobtrusively and transparently integrated in daily life and not becoming a source of potential problems and nuisance as often perceived today.

Emerging Global Risks: Digital systems evolve, but the threats and types of attacks also change continuously. There is, thus, a need to permanently survey and identify new attacks, to monitor potential network and service vul-

nerabilities and to look for new emerging risks. Examples of such risks at the global level include: excessive disclosure of private information, bullying, identity theft and squatting and predators masquerading.

International Test-beds and Datasets: Ways to interconnect test-beds should be explored, including the connection of test-beds that were developed as standalone by the different countries. Such test-beds would permit to share data sets and carry out validation in a co-ordinated manner, internationally. Cooperation can be established on various levels: means and results, approaches, infrastructures, software and data. The problem is how to federate such test-beds, taking into account cross-testing, mobility aspects and security policies as users move in and out of different environments. Additional problems to address for enabling the sharing and exchange of data and information relate to the intellectual property, interoperable data formats in repositories, the "diluted glory" factor, the confidentiality or the reproducibility of experiments. Two specific examples of future test-beds discussed were a test-bed on international application and software services that could be built on top of GENI with a number of application-level experiments and a test-bed for wireless or sensor networks.

Mechanisms for International Collaboration: Participants also discussed mechanisms that EU and US funding agencies could make available to assist the continued international collaboration. These included the establishment of a co-ordination type project that could act as catalyst to facilitate, drive and set up future activities in a systematic fashion. Ideas for additional international projects included dynamic service coalitions, privacy, and legal considerations for global cyber-security involving all relevant stakeholders (technologists from academia and industry, policy makers, legal, consumers). A number of realistic cooperation mechanisms based upon existing programmes in each country were also presented by representatives of the European Commission, NSF and DHS. The Japan Science and Technology Agency and Australian NICTA have indicated a willingness to run a parallel programme that could work alongside the EU and US efforts.

Links:

2nd EU-US Workshop on Cyber Trust (all presentations and position papers):
<http://www.iti.uiuc.edu/EU-US/>

Workshop report:
<http://www.securitytaskforce.eu>

1st EU-US workshop report on Cyber Trust:
http://www.securitytaskforce.org/images/stories/eu_us_cyber_summit_report.pdf

Please contact:

Jim Clarke, Waterford Institute of Technology
 E-mail: jarlarke@tssg.org

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

Science Foundation Ireland funded Online Dublin Computer Science Summer School

24 International Research Undergraduate Projects in Ireland for "Technologies for Aiding Human Memory".

ODCSSS the Online Dublin Computer Science Summer School is a paid research internship program for undergraduate students funded in part by the Science Foundation Ireland under their UREKA program. ODCSSS is a four-year collaborative internship program between the School of Computer Science and Informatics at the University College Dublin and the School of Computing at the Dublin City University.

The 2007 thematic focus of "Technologies for Aiding Human Memory" grew out of our experiences with the program in 2006. The 2006 students suggested the need for



ODCSSS student with TableTop Tissue Micro Array Visualisation.

undergraduates to see a social impact from their work, not just academic results. In 2006 there were 17 students in the program, which has grown to 24 students in 2007. This year we had over 100 applications from around the world. In 2007 we have students coming from Universities in Austria, Czech Republic, France, Germany, Hong Kong, India, Ireland, Slovenia, Spain, Thailand and the USA. The 2008 program opens for applications from 2nd and 3rd year undergraduate students in late 2007. Successful candidates are offered a tax-free scholarship of €300 per week (for 12 weeks) along with travel support.

The primary goal of this research program is to afford exceptional undergraduate students the opportunity to participate and contribute to exciting yet challenging research projects and to inspire them to go on to undertake research careers. Example projects in 2007 include, "Reminding Short-Term Memory Sufferers to Complete Routine Tasks" and "Using Multiple Sensors to Determine Posture". Past projects include "TableTop Tissue Micro Array Visualisation" and "Forensic speaker identification".

While focusing on world-class research, this internship provides the participants with exposure to a distributed, cooperative and collaborative research environment facilitated and supported by exchanges of the entire summer school cohort. Along with research seminars, research training, industrial lab visits, industrial research events, social events and research management seminars including intellectual property, patenting and commercialization.

Link:

<http://www.odcsss.ie>

Please contact:

Aaron Quigley, University College Dublin / IUA

Tel: +353 1 716 5357

E-mail: aquigley@ucd.ie

EuroSpaceward – Preserving and Improving Life on Earth by Going into Space

by Markus Klettner, Benoît Michel and Bradley Edwards

EuroSpaceward is a brand new European Association created by scientists believing in space exploration and settlement as a way to preserve the human civilization.

It does not require an extraordinary clairvoyant ability, in order to understand that our modern civilization can only survive and prosper on a long-term basis, if it succeeds to lastingly solve its inherent burning problems, like the irreparable destruction of natural habitats, the depletion of valuable raw materials of mother Earth, the ever increasing need for energy, or the potential devastating climatic changes induced by a polluting world economy.

This is where the European Spaceward Association (short: EuroSpaceward), a non-profit organization established in Luxembourg in March 2007 and a liaison-partner of the Spaceward Foundation in the United States, intends to make its contributions.


The non-profit association's main purpose is the promotion and dissemination of activities at European universities and European schools, in the public and in economy that are aiming at sensitizing to the need for the sustainable industrialization and settlement of space to preserve the planet Earth as well as the safeguarding of its biosphere and diversity of its species including the continuous improvement of the quality and safety of human life.

EuroSpaceward also engages in the introduction and coaching of development teams at European universities, technical and scientific centres aiming at the participation in the publicly accessible space programs that support the overall mission of facilitating the industrialization and the settlement of

space, like Elevator2010, NASA Beam Power and Tether Challenge, Mars Barn, Regolith Excavation Challenge, publicly open ESA projects, etc.

The flagship project of EuroSpaceward is focused on the Space Elevator; a concept for inexpensive and safe travel from Earth to space that only recently has been considered viable. By utilizing state-of-the-art technology, a ribbon can be strung between an anchor on Earth and a satellite beyond geosynchronous orbit. This ribbon will be ascended by mechanical climbers thereby providing ready access to space without using any rockets. Currently in the engineering stages this concept has attracted considerable interest internationally and has been the focus of international conferences and engineering competitions in the United States to promote development of the required technologies. With the assistance of EuroSpaceward researchers and engineers in Europe are now beginning to take active roles in this endeavour.

by courtesy of Spaceward foundation



The space elevator concept - climbers ascend a ribbon strung between an anchor on Earth and a counterweight in space.

Having been awarded with a grant by the Fond National de la Recherche of Luxembourg, EuroSpaceward is now in the position to organize the first European Space Elevator climber design workshop in cooperation with the University of Luxembourg and the Institute Gabriel Lippmann as well as with the Centre Spatial de Liège, the University of Liège and the Université Catholique de Louvain. The workshop is planned to last 3 days and will bring together high level researchers and engineers on Space Elevator Systems along with private industry to examine and develop designs for climbing systems and elevator tethers that can excel at NASA's Beam Power and Tether Challenge and move the technology development forward. The participants and session leaders are the world leaders in this work and the workshop should result in a strong advancement in climber and tether design.

Link:

<http://www.eurospaceward.org>

Please contact:

Markus Klettner, EuroSpaceward Association

E-mail: Markus.Klettner@alumnimail.isunet.edu

Introduction to the Special Theme

Service-Oriented Computing



by Jana Koehler and Gustavo Alonso

Service-oriented computing is an emerging cross-disciplinary paradigm for distributed computing, which is changing the way software applications are designed, delivered and consumed. At the heart of service-oriented computing are services that provide autonomous, platform-independent, computational elements that can be described, published, discovered, orchestrated and programmed using standard protocols to build networks of collaborating applications distributed within and across organizational boundaries.

Grid services and Web services are currently the most common forms of service for implementing service-oriented computing. Grid services provide the foundation for the distributed execution of long-running scientific computations over very large data sets using a standardized and stateful service interface. Web services provide the basis for the development and execution of business processes that are distributed over the network and available via standard interfaces and protocols. Technically, these two types of service have converged to a large extent in recent years, but the difference in application focus leads to a variety of complementing research questions.

The state of the art in service-oriented computing is characterized by the very early adoption of many forms of technology by the IT industry. In particular, in the area of Web services, large IT vendors are leading and driving many research activities. At the same time, many customers have begun to use Web services. However, it is a challenging task for many corporate IT organizations to migrate an enterprise towards a service-oriented architecture. Adopting service-oriented computing (or a service-oriented architecture as it is often called in the enterprise computing space) is a difficult endeavour given the fact that much of the required technology is at an early stage of maturity. Problems can be located in the usability of today's software tools for developing service-oriented applications, the incompleteness of the so-called Web services stack with its many standards, and the need to further mature many of those standards. Key standards are the Web Services Description Language (WSDL), the Business Process Execution Language (BPEL), and the Simple Object Access Protocol (SOAP). In addition, many WS-* standards (for example, WS-Addressing, WS-Policy, WS-Reliability, WS-Security) are addressing non-func-

tional aspects of Web services. Semantic Web standards such as RDF and OWL have been adopted by the first commercial tools and used to extend Web services with annotations that describe more precisely what functionality a service provides.

Besides this, numerous problems relating to user-friendliness and the maintenance of the implemented IT solution artefacts must be overcome. For example, the problem of ensuring the correctness and completeness of business process solutions using service-oriented computing principles has been only partially solved. Similarly, it is very difficult to quantitatively define what constitutes a good service-oriented architecture and how to determine the right granularity of the designed services in order to facilitate reuse. This is what service-oriented computing is about – the next evolution in the reuse of software combined with principles of loose coupling and distributed computations. Yet initial experience indicates that even Web services may not be as loosely coupled as originally expected, and there is a strong trend towards Representational State Transfer (REST) and REST-oriented development with very thin intermediate layers (in addition to WSDL/SOAP).

In Grid computing, services and service-oriented computing plays a big role, and to a certain extent it was in Grid computing that some of the early ideas about service-oriented architectures were first explored. This early adopter strategy led to diverging and independent specifications for Web services in the enterprise and in Grid computing. This phase is now over and the convergence between both views of services and the related infrastructure is almost complete.

However, Grid computing poses different challenges and requirements than do enterprise applications. For instance, the amount and type of data exchanged is quite different than in enterprise applications, and the dependencies between tasks in a composition are typically stronger since often there is not just a data or control flow dependency, but also a location dependency imposed by the cost of moving data. As another example, the Web Services Resource Framework (WSRF) is a specification tailored to Grid computing to enable closer control of stateful resources (eg computations) through Web services. This specification can be used in many applications but it addresses problems that are found primarily in computing grids rather than on enterprise applications. These and other differences are substan-

tial enough to require systems and approaches other than those encountered in enterprise computing, as the use of resources and the operations performed with and on services often involve a higher degree of control than conventional enterprise services.

Grid computing, however, still has to catch up with enterprise computing in terms of industrial strength tools and a better understanding of the development process. While a new form of software engineering is quickly evolving in the enterprise computing space, driven by the need to bring business and IT together, service composition and service architectures are still a matter of mostly ad-hoc development in Grid computing. Thus, what has been achieved so far is a leap forward over the state of the art from a few years ago. Still, much remains to be done to incorporate into Grid computing the discipline, techniques and lessons learned from software engineering. This is also true with regard to services, and will remain the main challenge in Grid computing for years to come.

The articles featured in this special edition reflect the recent trends in research on service-oriented computing, and address many of the aforementioned issues. Human collaboration is one of the strong drivers in the Grid and Web services area. It is therefore not surprising that non-functional aspects such as the quality of a service are at the heart of many research and development activities. Semantic technology and well-formalized analytical capabilities aim at improving tools to develop service-oriented applications. We are especially pleased to have several articles in this issue that illustrate the wide applicability of service-oriented computing including, for example, the shipping industry, customized shoe production, health care and satellite-based observation. This is a clear indication of the strengths and potential of service-oriented computing.

Please contact:

Gustavo Alonso
Computer Science Dept, ETH Zurich, Switzerland
E-mail: alonso@inf.ethz.ch

Jana Koehler
IBM Research GmbH, Zurich Research Laboratory,
Switzerland
E-mail: koe@zurich.ibm.com

The UK National Grid Service

by Katie Weeks


The National Grid Service (NGS) is the UK's Grid for academics from all fields. Our users range from chemists to bioinformaticians to social scientists. The ultimate aim of the National Grid Service is to provide a wide range of computing and data resources and services for all users without requiring them to know anything about Grid computing.

Having entered production in September 2004 with four core sites (at the University of Manchester, the University of Oxford, the Science and Technology Facilities Council Rutherford Appleton Laboratory and the White Rose Grid at Leeds University), the NGS is now well on the way to achieving this aim, and has grown to include six partner sites at Cardiff University, University of Bristol, Lancaster University, University of Westminster, Queen's University Belfast and the UK's supercomputer HPCx. The National e-Science Centre in Edinburgh is also an affiliate of the NGS.

More than five hundred users from all areas of research are currently registered on the NGS. Projects include large simulations of drug permeations through a membrane, studies of the molecular basis for HIV drug resistance, modelling of the human heart, computational modelling of aircraft structures and simulations of the generation of magnetic fields in the Sun and in planets. The NGS also provides an Oracle database service for the storage of large amounts of data, such as the database of possible crystal structures for the Control and Prediction of the Organic Solid State (CPOSS) project.

The NGS currently runs Globus Toolkit 2 on its core nodes. The second phase of NGS began in October 2006, and in the summer of 2007 will see each of the four core sites installing new clusters with Globus Toolkit 4, while retaining backwards compatibility with Globus Toolkit 2. Access is managed via X.509 digital certificates, which users must have before applying to the NGS. The move to Globus Toolkit 4 will enable Web service support and development. The NGS is working with users and developers in this area in order to better understand the requirements of such a service.

The diversity of users on the NGS requires a wide range of programs to be



With an account on the National Grid Service, researchers in the UK are able to perform simulations such as the pictured drugs permeating from water (represented by red and white lines) through a membrane (represented by turquoise lines). By using the NGS, the simulation time was cut from years per drug to about 2 weeks.

installed. As well as standard compilers such as Gnu, Java and Intel, many libraries and specific scientific packages are run on nodes within the NGS, including Siesta, Gaussian, ncbiBlast and dl_poly. Users are encouraged to request the installation of software useful to them, subject to licence constraints that affect the installation of many commercial products. There is also an area reserved on the STFC-RAL node for users to install their own and open-source software without making an official request.

Whilst the NGS has been in production for two years, it recognizes there is still a long way to go to meet all user requirements. The NGS support centre, led by the STFC Rutherford Appleton Laboratory, works to ensure that users are supported throughout their Grid career and that all requests for improvements are listened to. User forums are held twice a year to encourage users to

become involved in the development and future planning of the NGS.

Research and development for the NGS is always in progress. Some of the more recent research includes the installation and development of the EGEE glite Resource Broker to work with the Message Passing Interface (MPI) standard. The recently released Applications Repository, a portlet that allows users to submit jobs via a Web interface using standard job templates for software such as Gaussian and dl_poly, is undergoing continuous development in response to user feedback. And whilst the monitoring of individual usage has been running since the NGS entered production, the User Accounting System was only launched in October 2006 after months of work. The NGS Usage Accounting System was developed because no other Grid accounting system met the needs of the NGS. The system allows the monitoring and accounting of individual accounts, with automatic e-mail alerts when a user reaches 90% of their allocated quota, and automatic account deactivation when they exceed their quota. Further research and development is going into the Usage Accounting System to enable it to accommodate virtual organisation management (VOM) and storage services.

The NGS has grown steadily since entering production, but with plenty of room left and with more partner sites lining up to join and offer their resources, the future looks bright for the NGS. The user base is expanding as more people realize what the NGS can offer them, and with services and support developing all the time, this can only continue.

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

Please contact:
UK NGS Helpdesk
Tel: +44 1235 446822
E-mail: support@grid-support.ac.uk

Mobile Service Management in Service-Oriented Grids

by Tom Kirkham, Fredrik Solsvik and Robert Piotter

Akogrimo is an EU Framework 6 project developing an infrastructure to support the use of mobile services within Grid computing applications. This investigation focuses on the presentation of mobile services as valid resources for Grid applications to use. The project is nearing completion and has developed an architecture supporting eHealth, eLearning and Disaster and Crisis management testbeds.

What sets Akogrimo apart from many other service-oriented architectures (SOA) and mobile Grid applications is that Akogrimo focuses on the provision of applications that use resources sourced from mobile Grid services. This is in contrast to common methods of mobile Grid usage where the client is mobile and the services static: a one-to-many relationship. In the case of Akogrimo, both clients and services can be mobile.

The key advantage of SOA in Grid environments is that applications can be composed from multiple separate services that are executed to a specified workflow. This enables the SOA applications that run on Grid computing platforms to become more business-oriented and adaptive to user needs. Users also have increased flexibility in their choice of service provider, with service selection taking place close to execution time. Mobile Grid services add an extra dimension to this, as the workflow must adapt to service behaviour introduced with the addition of mobility.

In order to enhance this service composition and workflow management, Akogrimo has developed an infrastructure to support the selection and execution of

mobile services. This is based on both context management and service composition. The Akogrimo approach is to present service and network context monitoring to aid service composition based around four main architectural layers: the network, service providers, Operative Virtual Organization (OpVO) and Base Virtual Organization (BVO).

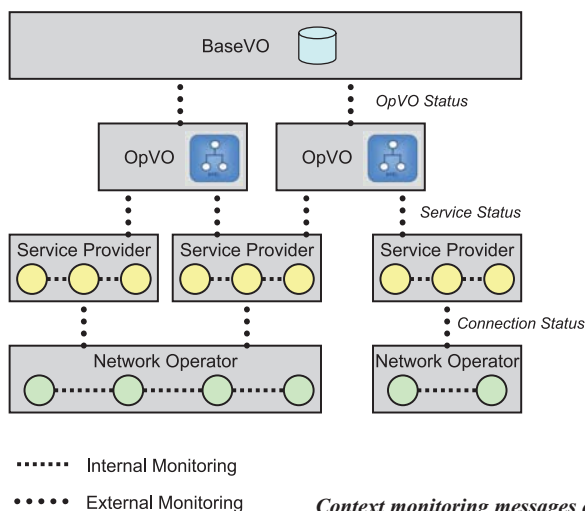
The Base VO is the head of the infrastructure and controls the behaviour of the services that pass information down to the bottom two layers. The Operative VO layer allows the BVO control to be abstracted and applied to specific applications, the design being an Operative VO per application instance. This allows specific services to be grouped together in environments that are focused on a specific application delivery. The service providers are mobile and linked to the network, and data related to mobility is monitored by services in the OpVO within specific context management services.

There are three testbeds. One (eLearning) is mainly used as a case study for developing individual services using the Akogrimo infrastructure. The eHealth testbed uses a service-oriented workflow that invokes application serv-

ices based on the infrastructure. This twelve-partner testbed was demonstrated forty times at a major European conference (IST2006). The DHCM (Disaster Handling and Crisis Management) testbed is being developed at present and the infrastructure and partnership are being extended for this.

The key innovations in these testbeds have been in the services that enable the Akogrimo applications to traverse the levels of the mobile Grid infrastructure. This has been achieved by services that represent network and Grid middleware services. These services work by processing subscriptions and monitoring changes in mobility of the services at the network level that may affect the larger workflow of the application at middleware level.

As Akogrimo goes into the future the focus is on increasing the level of communication between these layers using approaches that are increasingly available in similar fields such as the Semantic Grid. In addition, it is intended that a wider range of mobile positioning devices be used, introducing the concept of mappings of space to the applications. A bid is being developed to explore this in more detail.



Link:
<http://www.mobilegrids.org/>

Please contact:
Julian Gallop, STFC, UK
E-mail: j.r.gallop@rl.ac.uk

Tom Kirkham, STFC, UK
E-mail: t.d.kirkham@rl.ac.uk

Fredrik.Solsvik, Telenor, Norway
E-mail: Fredrik.Solsvik@telenor.com

Robert Piotter, HLRS, University of
Stuttgart, Germany
E-mail: piotter@hlrs.de

A4C Support for Commercialization of Next-Generation Grid Services

by Hasan, Peter Racz, Cristian Morariu, David Hausheer and Burkhard Stiller

The development of Grid technology has reached a point at which it is suitable for commercial deployment in a multi-provider environment. The EU Project Akogrimo is driving this development toward the support of mobile participants of a Virtual Organization employing Next-Generation Grid infrastructure. With the support of A4C functions developed within Akogrimo, such mobile Grids can now be commercially exploited. These A4C functions are the focus of this overview, and comprise user authentication, resource access authorization, multi-domain accounting of resource usage, auditing of compliance with SLA, and charging of resource consumption.

Grid technology is evolving from a niche market, in which it solely addressed the management of shared and distributed resources, into a framework which incorporates knowledge-related and semantics-driven Web services allowing for applications in a broad business context. This evolution has led to the so-called Next-Generation Grid (NGG). In addition to traditional Grid applications, envisioned NGG applications include e-business, e-health, e-government, and e-learning.

Until recently however, the Grid community has not so far considered the issue of mobility. Given the vast growth in the number of mobile Internet users, the EU Project Akogrimo is therefore aiming to advance the pervasiveness of Grid computing across Europe by uniting concepts and results gained in the systems beyond 3G (B3G) and the Grid community. Driven by business requirements, authentication, authorization, accounting, auditing and charging (hence the name A4C) are crucial functions for the commercial deployment of NGG technology for providers and users of mobile Internet Grid services. In order to present mobile users with a transparent view of the use of orchestrated services, A4C functions must consider interactions between providers.

Next-Generation Grid and Mobile Dynamic Virtual Organization

A Virtual Organization (VO) is understood to be a temporary or permanent coalition of geographically dispersed individuals, groups, organizational units or entire organizations that pool resources, capabilities, and information to achieve common objectives. A VO can provide services, including high-level resources such as knowledge (eg user and device context information and state information of VO components),

and it may comprise various types of service providers, such as Grid service providers, content providers, network service providers, and even other VOs. A Mobile Dynamic VO (MDVO) is composed of potentially mobile participants where contracts are dynamically established. An NGG allows for the formation of MDVOs to solve complex problems across different network technologies and provider domains. Figure 1 depicts the NGG organization model in which various resources and

Akogrimo NGG architecture whose components are grouped into three layers: Grid Application Support Services (GASS), Grid Infrastructure Services (GIS), and Network Middleware (NM). GASS components are responsible for Service Level Agreement (SLA) negotiations, establishment of VOs including operational VOs, and creation of workflows. GIS components manage and monitor the execution of tasks as specified by GASS components, measure resource usage and service perform-

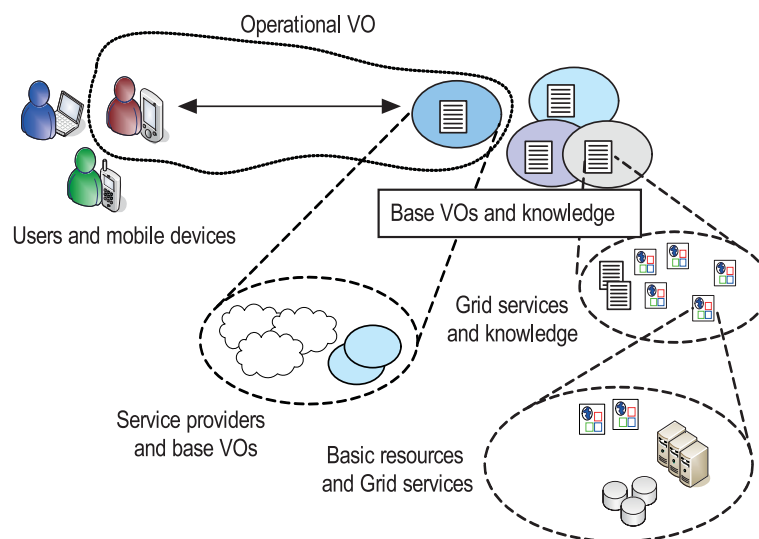


Figure 1: NGG Organization Model.

actors are brought into relation. The base VO comprises pools of potential resources, services and providers that are combined into an instantiation of the VO for one user or customer. This is termed an operational VO.

A4C Integration into Next-Generation Grid Infrastructures

The integration of A4C into an NGG infrastructure is the key to commercializing Grid resources, services and knowledge that are distributed across multiple domains. Figure 2 depicts the

ance, and monitor compliance with the respective SLA. Finally, NM components provide for network QoS support, functionality to support user-, device- and session-mobility, management of contexts and service discovery, and A4C functionality to allow for a controlled access to and an accounted usage of resources.

Two key concepts form the basis of authentication in Akogrimo: Single Sign-On (SSO) and Anonymity. SSO is achieved by integrating the Secure

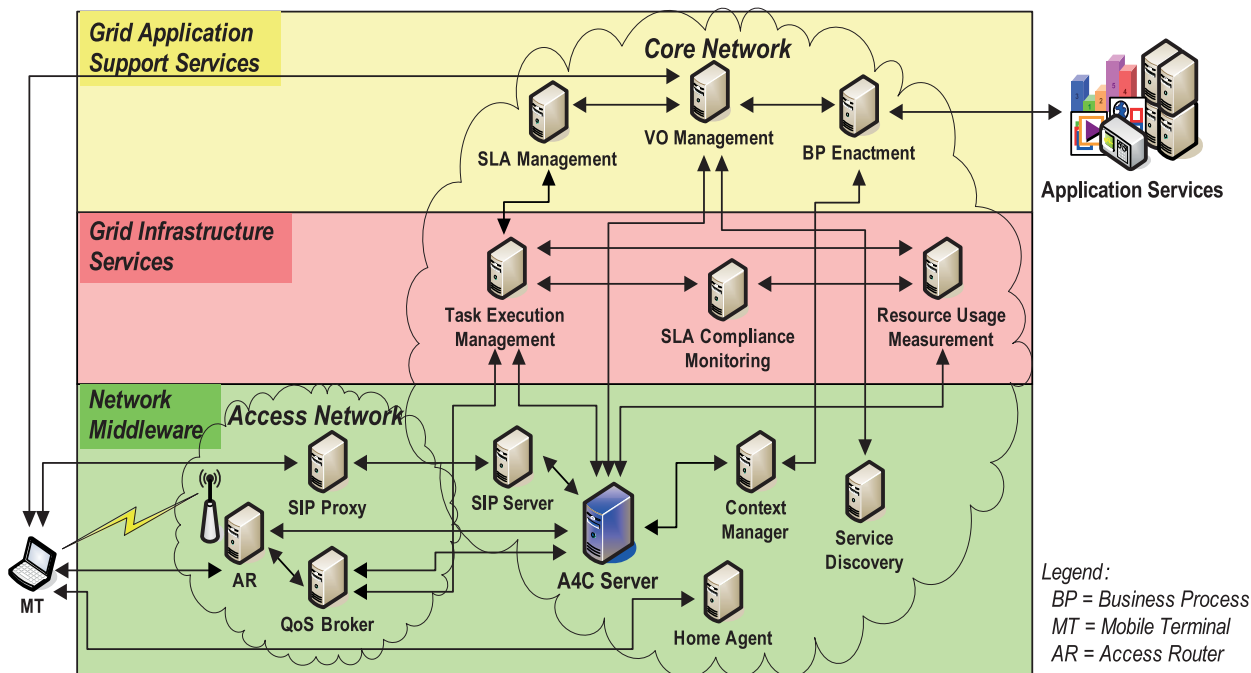


Figure 2: Akogrimo A4C in the Next-Generation Grid Architecture.

Assertion Markup Language (SAML) into the A4C infrastructure. Anonymity is supported at two different levels: pseudo and full anonymity. Under pseudo-anonymity a user retains the same virtual identity (a pseudonym) for each service request and access. Under full anonymity a new virtual user identity is created each time a user requests a service from a service provider.

With respect to authorization, two levels are distinguished within Akogrimo. At the first level, network access authorization is performed by the access router and the QoS broker. After a successful authentication, users may only use service bundles to which they have subscribed. Service bundles are stored by the A4C server and are provided to the QoS broker upon request. The second level of authorization lies in Grid service-related layers. The authorization at this level enables access to Grid services to be controlled according to existing policies in the operational VO as well as in the administrative domain of participating service providers.

The A4C infrastructure supports accounting for resource and service usage in a mobile, multi-domain service-provisioning environment. Accounting data related to network and Grid resource usage are gathered on the vari-

ous network components that provide the service, and then collected and stored by the A4C server in the form of accounting records. To correlate accounting records originating from different components, accounting sessions are defined. Accounting sessions are related to the service usage of a user and bind accounting records together. The session model enables a hierarchical session structure in which services might use other services: this is highly probable in a Grid environment. The A4C server also receives and stores auditing events that express the degree of SLA compliance during service consumption.

Furthermore, the A4C server performs charging by applying a charging scheme to collected accounting records and auditing events. The charging scheme describes the charge calculation rules and tariffs to be applied and is specified by an XML-based representation. According to auditing events a provider can apply penalties or discounts to the final charge. Finally, the user's home provider prepares a single bill for all accessed services.

Interested readers can obtain more detailed information from documents of the Akogrimo project as given on the links. Finally, we would like to thank Martin Waldburger for his contribution in technical discussions.

Links:

Akogrimo Project:
<http://www.akogrimo.org/>

Final Integrated Services Design and Implementation Report:
<http://www.mobilegrids.org/modules.php?name=UpDownload&req=getit&lid=110>

Building Grids for Europe:
http://ec.europa.eu/information_society/policy/nextweb/grid/index_en.htm

Please contact:

Dr. Hasan
 University of Zurich, CSG@IFI,
 Switzerland
 Tel: +41 44 635 75 86
 E-mail: hasan@ifi.uzh.ch

Collective Completion: How do YOU Complete your Unknown Colleagues?

by Magnus Boman, Kristofer Franzén and Fredrik Espinoza

Researchers at SICS are developing new middleware for integrating competence and knowledge management tools to make knowledge work more efficiently while leveraging individual interests and social networks.

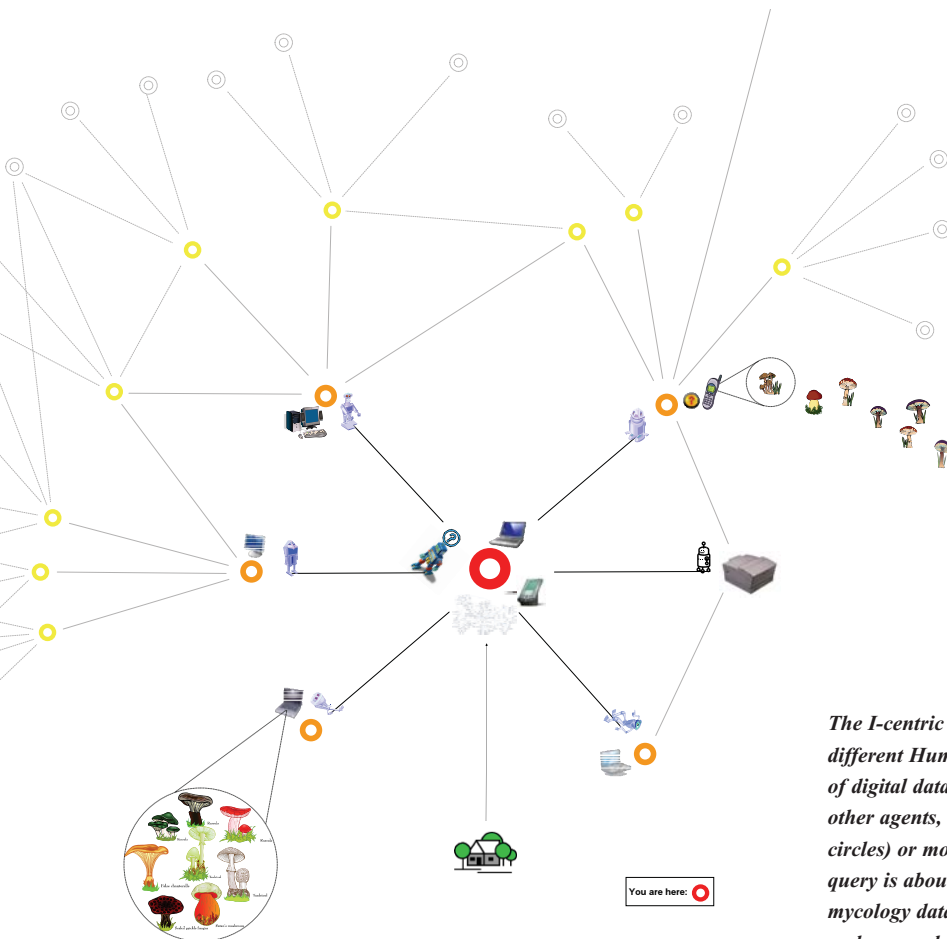
Successful cooperation depends not only on the ability to easily find, use, share and develop knowledge and competence, but also on participants' interests, experiences and skills. Finding, exchanging, aggregating and relaying knowledge are actions we all perform on a daily basis as part of cooperation. People and organizations also often find it necessary to find, contract, reimburse (financially) and evaluate people with specialist competence. The fact that people who cooperate and collaborate generally have common interests is usually understated, but such common interests strongly impact on the results of the cooperation. A network of people of varying skills and experience is a valuable source of help in solving concrete problems, but describing the skills necessary to do this is often a difficult task.

When it comes to ICT tools that support the search for skills and competences and the features of collaboration mentioned above, people use a plethora of digital support systems, and will probably continue to do so for the next five to ten years. Although each tool is well suited to its specific usage, there is a lack of support for integrating and aggregating solutions as well as for personalizing them to individual needs. The social networks that connect people are not efficiently exploited, partly because relativizing and centring a social network to a particular user is a complex problem.

We propose to improve the situation by introducing Human Grid between the human users and the underlying support tools. This is a layer for integrating services and connecting skilled people

in a user-centred network. It is meant not to increase the problem of disparate systems, but rather to enhance the user's desktop with an organizationally and individually adapted view which is easily turned on or off, and which includes the ability to run the usual tools from the usual providers. To achieve this increased utility without imposing yet more intrusiveness on the users' work processes or physical work tools requires advanced solutions to many separate problems that are rarely dealt with collectively.

Physicists, medical researchers, engineers and other work groups already use Grid computing: it is an effective way of utilizing the computing power of distributed and networked computers when and where it is most needed. In Human Grid the resource being tapped



The I-centric view of Human Grid where the user at the red circle has two different Human Grid-enabled devices, access to competence in the form of digital data, and a personalised agent. This agent sends a query to the other agents, each representing another user at social distance 1 (orange circles) or more (yellow circles), over a P2P network. This particular query is about mushrooms, and is eventually answered by accessing a mycology database (bottom left), feedback from the user context (bottom), and a snapshot taken by a camera phone (top right).

is not computing power but human competence. The two paradigms agree in that the use of the network resources should have a minimal impact on the user and the required infrastructure should already be in place. What is missing is the middleware that allows the user to switch on the network view of regular activities. The primary concern for such middleware must be the ease of use, transparency and scalable accessibility for its users. It should also be secure and should strive to protect users' privacy, for example by making available the proposed uses for information, competence and knowledge that may be harvested in a particular situa-

tion. Trust mechanisms, similar to those used to preserve and share source code and intellectual property (eg Creative Commons) should be included. It should be possible to convert implicit competence, on group, project and organization levels, to explicit competence when the demand arises.

We are working on a four- to five-year time horizon, with the first phase of the project now behind us. Work is in progress in a number of areas: social networks for tying together users; language technology for gathering the necessary traces of human activity and competence; and service-oriented architectures

for analysing the relationships between underlying services and solutions for their integration and aggregation.

The first industry partner of the project is HP Labs, Palo Alto, which has demonstrated a strong interest and prior research in the area. The primary researchers at SICS are Magnus Boman, Kristofer Franzén and Fredrik Espinoza.

Please contact:

Magnus Boman, Kristofer Franzén,
SICS, Sweden

Tel: +46 8 633 1500

E-mail: mab@sics.se, franzen@sics.se

Workflow Management Systems for Grid Computing

by Carmen Bratosin and Wil van der Aalst

The Architecture for Information Systems group of the Technische Universiteit Eindhoven (TU/e) in the Netherlands has built up extensive knowledge in the field of Workflow Management Systems and Process Mining. Since 2006, the group has begun to apply this knowledge in a new and dynamic research area: Grid computing. Four research perspectives are currently under investigation.

Software systems are becoming increasingly complex. To cope with this, systems are often divided into a number of autonomous components whose work is coordinated: this coordination of components and services represents one of the main challenges in software engineering. Two important application fields of coordination are Grid computing and workflow management. Grid computing is mostly used in computational science while workflow management is used for business applications: we try to bridge the gap between these two areas in order to make further progress in both of them. Over the last decade we have gathered a great deal of experience in process modelling, analysis and enactment. Our workflow patterns have become a standard way to evaluate languages and the workflow management system YAWL (Yet Another Workflow Language) is one of the most expressive and mature open-source workflow systems available today. Moreover, we specialize in process analysis. Using Petri nets as a theoretical foundation, we have been able to analyse a variety of real-life process models ranging from BPEL (Business Process Execution Language) and workflow specifications to

the entire SAP reference model. In recent years, we have focused on the analysis of processes based on system logs. The ProM framework developed at TU/e provides a versatile toolset for process mining, which seems to be particularly useful in a Grid environment.

Until now, the Grid computing community has focused primarily on infrastructure. Grid software has been designed that allows users to submit their 'problems' to the Grid. Less work has been done on how to model such problems efficiently. In addition, most applications place the correctness properties in the hands of the user.

We are therefore applying our knowledge of Petri-net modelling and analysis, workflow patterns, process mining and concrete workflow technology to Grids. This involves research in the following areas:

- *Grid modelling:*
Many definitions of Grids exist, and in many cases technological aspects and hyped terms hide the essence of Grids. We use a mixture of Petri nets and UML modelling to build formal/con-

ceptual models for Grid computing. Here we emphasize the link between the distributed nature of Grids (where resources play an important role) and workflow processes. The main purpose is to formalize the concept of a Grid and to fix a particular interpretation while highlighting interesting research questions.

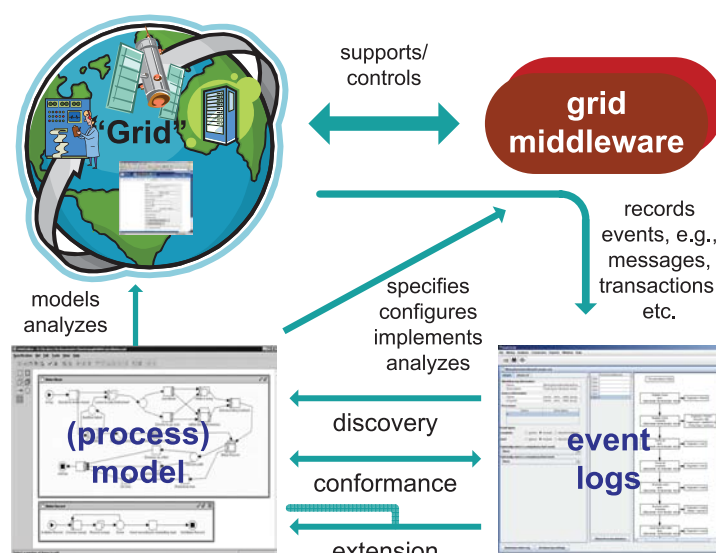
- *Analysing Grid models:*
Using techniques based on Petri nets, we analyse different mechanisms used in Grid workflows, with the goal being to transfer correctness notions such as soundness to them. We also try to find new properties based on the specific Grid behaviour (eg multiple instances of the same process, resource allocation, and distributed management).
- *Analysing Grid logs:*
In a Grid environment many events are logged and the performance of the system is of the utmost importance. The application of process-mining techniques is therefore of interest, to assist in the configuration of Grids and the on-the-fly optimization of processes.

- *Building a process-aware Grid infrastructure:*

Using a combination of Globus, YAWL and ProM we want to realize a more 'process-aware' Grid. By linking a fundamental enabling technology for the Grids (Globus) to a powerful process engine (YAWL) and state-of-the-art analysis tools (ProM), we obtain an interesting environment for experimentation.

The figure illustrates the scope of the project. On the one hand, we analyse Grids by modelling them in terms of Petri nets. Similar models are used for the configuration of the process perspective of Grid middleware (in our case a mixture of Globus and YAWL). On the other hand, we collect event logs via the middleware layer and use these for process mining, process discovery (automatically deriving models by observing the Grid), conformance checking (to check whether 'the Grid' is behaving as expected) and model extension (eg to project performance indicators onto a process model).

All of the aspects shown in the figure have been extensively investigated in the context of workflow management systems and service-oriented architectures using BPEL engines. For example, we have been doing conformance testing in the context of Oracle BPEL, and process discovery and process verification in the



Analysis and mining.

context of IBM WebSphere. We have also evaluated many process engines using the so-called workflow patterns and provided semantics and analysis techniques for process-modelling languages ranging from BPEL and YAWL to BPMN and EPCs. The next step is to apply this in a Grid environment using both Globus and YAWL.

The research is supported by the Netherlands Organization for Scientific Research (NWO) in the context of the project Workflow Management for Large Parallel and Distributed Applications. TU/e is participating in this project with the group of Professor Farhad

Arbab of CWI. The project started in 2006 and its duration is four years.

Links:

<http://www.workflowpatterns.com>
<http://www.processmining.org>
<http://www.yawl-system.com>
<http://www.win.tue.nl/ais>
<http://www.globus.org>
http://www.nwo.nl/nwohome.nsf/pages/NWOA_6R7FYJ

Please contact:

Carmen Bratosin, Eindhoven University of Technology, the Netherlands
 Tel.: +31 40 247 5144
 E-mail: c.t.bratosin@tue.nl

Towards a Commercial IT Service Delivery

by Christophe Ponsard, Gautier Dallons, Stéphane Mouton and Philippe Massonet

Many companies are currently moving or considering the move to a service-oriented architecture (SOA) model. This model holds the promise of flexibility and cost reduction by enabling business-level integration across organization boundaries. While this supports large companies by allowing them to better structure both internally and externally, it is also of benefit to small and medium enterprises (SMEs), as their highly value-added expertise can be made available more easily and with fewer overheads. However, achieving this vision still requires a number of challenges to be solved: management of quality of service, accounting, security and trust, interoperability and so on. CETIC, a Belgian research and technology transfer centre, is taking an active part in this work through a number of projects at both European and regional levels. Here we describe how the above challenges are presently being addressed, both at research level and within real-world deployments occurring in connection with that research.

Managed Quality of Service

Achieving commercial quality means being able to provide guarantees about quality of service. This is not an easy task given that service delivery can require the dynamic composition of a

number of elementary services in an open environment. Each step in such a chain is a potential point of degradation or failure. With success comes a strong federation of companies, each providing an improved return to all of them. Such

a federation has been achieved in the Belgian 'Walloon Marshall' plan, through the Walloon World Wide Space Application (3WSA) project. The coordinator of this project is Spacebel, an SME active in the space domain. The

developed platform is derived from the ESA Service Support Environment (<http://services.eoportal.org/>) and aims at chaining commercial services from various providers and processors of space data, each of which is supported by specific SMEs. In line with this effort, the following research issues are being transferred:

- Quality assurance, possibly through a defined certification process, for checking the adherence of a service implementation to its specification. This means service evaluation prior to deployment (for functional aspects) and at run-time (for non-functional aspects).
- High-availability assurance through a number of failure-provisioning mechanisms at platform level. In resource-

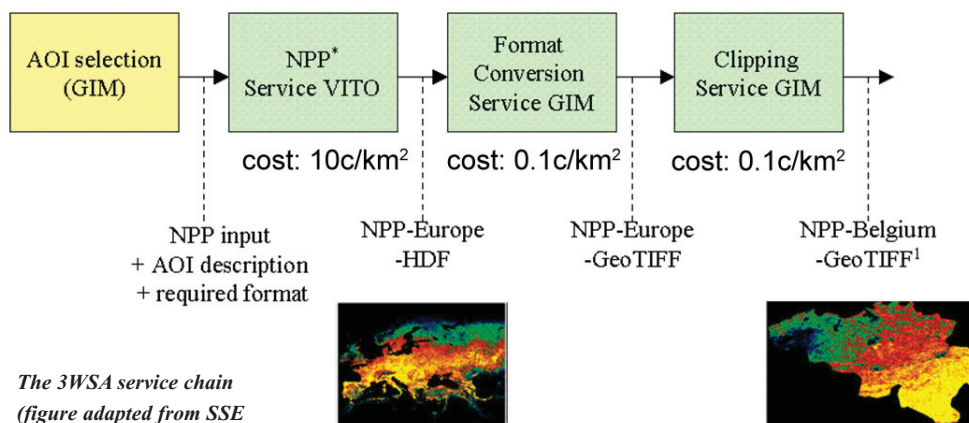
Accounting

New delivery models consider software as a service (SaaS), and move away from a licensing model. On top of this, billing may require more elaborate models if the final result has been produced by the chaining of a number of services, each of which brings some value for the customer in its business domain. This added value can be used by the accounting model in addition to technical costs such as CPU, RAM or disk usage. Platform infrastructure must therefore provide a means of tracking contributions in service chains, aggregating customer usage, applying pricing schemes (eg per usage, per time frame), and allowing data injection into accounting applications to finally bill customers. Such mechanisms are currently being investigated within the

work is currently being investigated in the GridTrust FP6 project. It is also being applied in some of the eighteen Grid business experiments of the BEin-GRID FP6 integrated project and in the previously described 3WSA project.

Perspectives

SOA is now past the early adopter phase. The mainstreaming process requires the above challenges to be addressed but also the emergence of a unifying framework. The NESSI European platform, which has gathered key industrial players and a strong research community, has the ambition to deliver this. NESSI is structured around problem- and domain-oriented workgroups, a number of which CETIC is involved in as they relate to the above challenges. The European Commission is also



The 3WSA service chain (figure adapted from SSE white paper, ESA 2004).

intensive platforms such as Grid computing, this can be achieved through techniques such as redundant executions, failure anticipation and checkpointing. Such work is currently being finalized by the HPC4U project, led by the University of Paderborn.

- Contract with failure clause and risk assessment, where a service contract between provider and customers specifies agreed service levels. Failure and associated penalties can also be part of this. Reliability figures can be gathered and published by independent third parties (brokers) who typically also carry out negotiations. A practical risk assessment approach can then be applied to reach the optimal price/penalties. Such an approach, based on WS-Agreement protocol extensions, is currently being developed by the AssessGrid consortium, led by the University of Berlin (TUB).

3WSA platform mentioned earlier, with each data/processing service being accounted for based on the defined area of interest as shown in the figure.

Security and Trust

Service orientation also means relying to some extent on externalized infrastructure, but security and trust issues such as data confidentiality, service availability and trusted infrastructure represent major barriers. To cope with these, new concepts such as Virtual Organization and service federation are being developed. These rely on the design, deployment and management of an adequate set of policies through which it is possible to enforce authentication, access control and confidentiality in a heterogeneous, dynamic and open environment. This does not mean enforcing strong security rules everywhere, but only where they are needed (ie based on actual trust relations). Such

strongly supporting this in the 7th ICT framework programme. The massive submission to objective 1.2 shows how important SOA is considered to be for achieving strong ICT support for the European economy of the 21st century. This is particularly the case for the federation of SMEs – the major workforce in Europe – around high-value services.

Links:

- <http://www.cetic.be>
- <http://www.assessgrid.eu>
- <http://www.gridtrust.eu>
- <http://www.beingrid.eu>
- <http://3wsa.cetic.be>
- <http://www.nessi-europe.com>

Please contact:

Christophe Ponsard
CETIC Research Centre, Belgium
Tel: +32 71 91 98 26
E-mail: cp@cetic.be

Reliable and Inexpensive QoS Monitoring in Service Markets

by Radu Jurca, Walter Binder, Boi Faltings

We propose a novel infrastructure for Quality of Service (QoS) monitoring that promises a significant reduction in monitoring costs. Our approach uses an incentive-compatible reputation mechanism to accurately estimate QoS by aggregating quality ratings from the clients.

Managing change is becoming increasingly important in a dynamic and quickly evolving economy. Consequently, organizations need flexible IT systems that can easily support new business processes and strategic objectives. Service-oriented computing enables the construction of such applications by orchestrating a variety of services that offer basic functionality. Recent efforts have materialized into standards and tools that significantly facilitate the construction and interoperability of services.

Since services are often developed and run by different organizations, they are generally provided under a contract (or Service Level Agreement) that fixes both the type and quality of the service to be provided, and the penalties to be imposed if these are not met. For example, a payment service that is part of an online shopping site may incur certain penalties if it fails frequently. In another example, the provider of a communication service may be penalized if they fail

to meet their guaranteed bandwidth. Yet another example is the provisioning of computation or data management services through a network: if a certain efficiency and capacity is guaranteed in an agreement, penalties may be incurred if these services are not provided.

An essential requirement for such service provisioning is the ability to monitor the QoS that is actually delivered. As the monetary value of individual services decreases, the cost of providing accurate monitoring takes up an increasing share of the cost of providing the service itself. For example, with current technology, to reliably monitor the quality of a communication service requires constant communication with a neutral third party and would be almost as costly as providing the service itself. This problem remains a major obstacle to the wider adoption of a service-oriented economy.

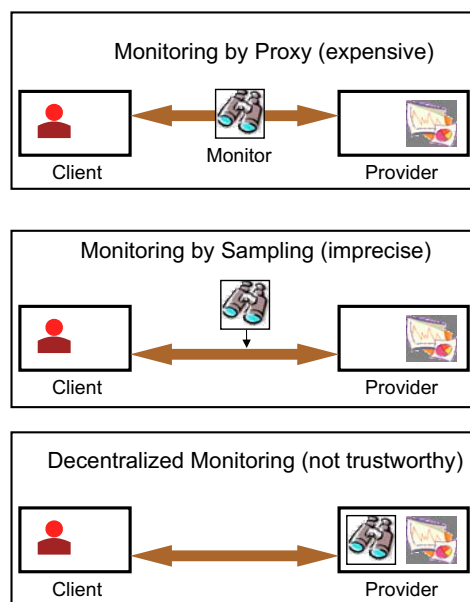
As a joint project between the Artificial Intelligence Lab of Ecole Polytech-

nique Fédérale de Lausanne (EPFL) and the University of Lugano, we are developing alternative mechanisms for monitoring QoS. The basic idea in our approach is to estimate the quality delivered by a service, based on the feedback provided by clients. Clients run the monitoring code, and periodically report feedback to a reputation mechanism (RM). The RM aggregates the reports and outputs QoS estimates for each service. The advantages of this approach are that (i) the RM acquires information on most transactions without actually being a bottleneck (there are no real-time constraints for reporting feedback, and the result of several interactions may be compressed into one feedback message), (ii) the monitoring process is as precise as possible (an immediate consequence of the first point), and (iii) service providers cannot directly tamper with the monitoring process.

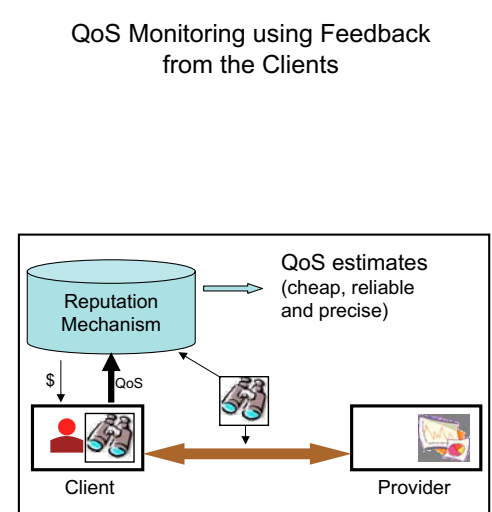
A practical mechanism must, however, ensure that clients report feedback honestly. Consider for example a client who

Traditional approaches to QoS monitoring involve: (i) centralized monitors that proxy every interaction, (ii) centralized monitors that sample the interactions or (iii) decentralized monitors. In our approach, most reports come from clients. A small percentage of interactions may also be sampled by a specialized monitor. Clients are paid in such a way that it is in their best interest to report honestly.

Current Approaches to QoS Monitoring



QoS Monitoring using Feedback from the Clients



knows that by reporting negatively, the resulting QoS estimate might entitle him to collect penalties from the service provider. The client can tamper with the monitoring code to make the provider pay penalties and thereby obtain a share of these. The novelty of our approach is to use economic incentives rather than hard security to obtain reliable information. Clients get paid for reporting feedback, and the amount depends on the values of other reports submitted by other clients. These feedback payments can be designed such that it is provably best for every client to report the truth. In game-theoretic terms, honest reporting becomes a Nash equilibrium, where any individual lie decreases the expected payment of the reporter. Misreporting is thus uninteresting rather than impossible.

Initial experiments have provided very encouraging results. We have implemented a QoS monitoring framework

based entirely on existing standards and technology. The extensions required for the collection and payment of feedback are lightweight, and do not introduce significant overhead. The payments for feedback can be made as low as one percent of the cost of the service, and the overall cost of monitoring can be decreased by an order of magnitude as compared to traditional techniques.

In our ongoing research, we are exploring an architecture that integrates reputation mechanisms with service directories and a publish/subscribe infrastructure. Clients subscribe for types of services in which they are interested and are notified upon significant QoS changes of certain services or upon the availability of matching services. Clients use this infrastructure to continuously evolve and improve their service-oriented applications. We will continue to use game-theoretic tools to ensure that the information distributed by the repu-

tation mechanism is reliable, and to ensure that self-interested parties have a rational incentive to adopt the behaviour recommended by our mechanisms.

Links:

http://dx.doi.org/10.1007/11596141_30
<http://www.2007.org/paper420.php>

Please contact:

Radu Jurca,
EPFL / SARIT, Switzerland
Tel: +41 21 693 6679
E-mail: radu.jurca@epfl.ch

Walter Binder
University of Lugano, Switzerland
Tel: +41 58 666 4303
E-mail: walter.binder@unisi.ch

Boi Faltings
EPFL / SARIT, Switzerland
Tel: +41 21 693 2738
E-mail: boi.faltings@epfl.ch

WS-DIAMOND: An Approach to Web Services – Monitoring and Diagnosis

by Luca Console and Mariagrazia Fugini

The WS-DIAMOND project aims at developing a framework for Web services that are endowed with self-diagnosis and self-repair capabilities. There are two main goals: the definition of an operational framework for self-healing service execution of conversationally complex Web services, where monitoring, detection and diagnosis of anomalous situations are carried out and repair/reconfiguration is performed; and the definition of a methodology and tools for service design that guarantee run-time diagnosability and reparability.

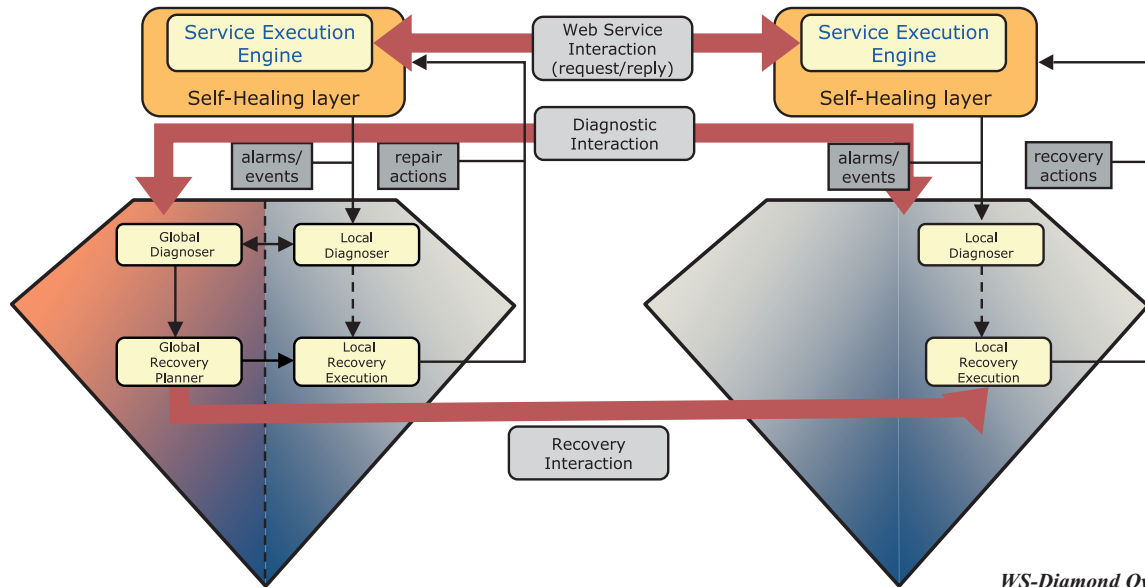
The WS-Diamond Project, funded by the EU commission under the FET Open Framework (EU IST FET-STREP n.516933), is developing a framework for self-healing Web services; that is, services able to self-monitor, to self-diagnose the causes of a failure, and to recover from both functional failures (eg the inability to provide a given service) and non-functional failures (eg loss of Quality of Service – QoS). The focus of WS-Diamond is on composite and conversationally complex Web services. The second goal of WS-Diamond is to devise guidelines and tools for designing services in such a way that they can be easily diagnosed and recovered at execution time, as well as tools to support the design of complex self-healing processes.

WS-Diamond commenced in September 2005 and will end in February 2008. The project involves universities and research institutions in Italy (Università di Torino – the coordinator – and Politecnico di Milano), Austria (University of Klagenfurt and University of Vienna), France (LAAS-CNRS-Toulouse, IRISA-Université Rennes, Université Paris Sud), and the Netherlands (University of Amsterdam).

The project assumes that the availability and reliability of complex services will be of paramount importance in the near future. Indeed the reliability and availability of software, together with the possibility of creating self-healing software, is recognized as one of the major challenges for IST research in coming years. Hence, WS-Diamond

research concerns a number of ‘grand challenges’ as described within the Service-Oriented Computing research roadmap at all levels: dynamic connectivity capabilities based on service discovery, QoS-aware service composition, and design principles for self-healing.

This work considers complex Web services, described using Web service workflow languages and frameworks such as BPEL4WS and extended Petri Nets. These services include mechanisms for augmenting processes with monitoring process functionalities using a methodological approach that focuses on exception handling and compensation mechanisms. Methodologies and tools will achieve adaptive Web-based process execution based on



WS-Diamond Overall Architecture.

flexible services. Attention is paid to conform to interaction patterns between organizations and to provide inherent flexibility and fault tolerance in process execution.

In its first phase (September 2005-January 2007), WS-Diamond designed and developed a platform for self-healing execution of complex Web Services, concentrating on run-time faults; the design issues are faced in the second phase of the project. This has led us to define the types of faults that can occur and that we want to diagnose, namely:

- functional faults and specifically semantic data errors (eg wrong data exchanges, wrong data in databases, wrong inputs from user)
- QoS faults.

In this first phase, WS-Diamond concentrated on orchestrated services, even though some proposed solutions already take choreographed services into account. We extended Web service execution environments to include features useful in the support of the diagnostic/fault recovery process. An architecture that supports self-healing service execution was then defined. The architecture provides support for associating a diagnostic service with each application service, for gathering observations about service execution (eg data exchanged between services) and pro-

vides a repair service as sets of recovery and repair actions. The architecture also includes a monitoring service aimed at identifying QoS problems, and a repair plan generator and executor that support the execution of recovery plans on the basis of the diagnostic information.

We defined a catalogue of faults and possible observations, and proposed an architecture for the surveillance platform. The correctness of the distributed diagnosis algorithms has been proved formally. Repair is based on repair actions (retry, compensate, substitute etc) and plans (generated online or pre-prepared offline for a given process) which are executed if a failure occurs in the process and a fault is diagnosed. Repair is thus characterized as a planning problem, where the goal is to build the plan of recovery actions to be performed at run time in order to recover from service and application errors.

Future activities will work on the complete definition of a method for designing self-healing Web services, focusing on the detailed design and development of a platform for observing a set of symptoms occurring in complex applications, for executing the distributed diagnosis, and for generating, selecting, and executing minimal repair plans (eg with respect to cost functions).

Link:

<http://wsdiamond.di.unito.it/>

Please contact:

Mariagrazia Fugini
Politecnico di Milano, Italy
Tel: +39 02 23993405
E-mail: fugini@elet.polimi.it

Formal Modelling and Verification in Service-Oriented Computing

by Maurice ter Beek, Stefania Gnesi, Fabio Martinelli, Franco Mazzanti and Marinella Petrocchi

Formal methods and tools are a popular means of analysing the correctness properties of computer network protocols, such as safety, liveness and security. First the protocol under scrutiny is described in a formal language, which often results in a more precise definition of its function. Subsequently, the properties to be analysed are specified in a suitable logic. Finally, to decide whether or not the protocol fulfils certain properties, automatic tools are used to analyse it. The outcome either proves the protocol to be correct with respect to the relevant properties or shows there to be a problem.

Recent trends in telecommunication for the delivery of services such as Web applications have highlighted the need to improve and extend the current network protocols. In this context, we have recently been involved, in collaboration with Telecom Italia Labs, in the formal modelling and verification of some protocols for service-oriented computing. This work is part of the EU-funded research project SENSORIA (FP6-2004-IST-FETPI). Its aim is to develop a novel comprehensive approach to the engineering of software systems for service-oriented overlay computers.

Here we briefly describe our experience in evaluating the suitability of the service-oriented application approach for the delivery of telecommunication services, and in investigating possible improvements and extensions. In particular, we present the formal analysis of a protocol for identity management (end-user identity in a federated context; application versus end-user identity), and a protocol definition supported by formal analysis. Readers are encouraged to refer to the website of the SENSORIA project for more details.

Identity Federation Protocols for Web Services

Telecom Italia recently proposed a network protocol that permits end users to access services through different access networks (both mobile and fixed) without explicitly providing any credentials, while the service application level can nevertheless trust the identity and authentication information provided by the access networks. As a result, service providers (SPs) identify a user with the authentication procedure performed by the network provider. After identity federation, a single sign-on suffices for a user to access all services belonging to the same circle of trust of SPs, while

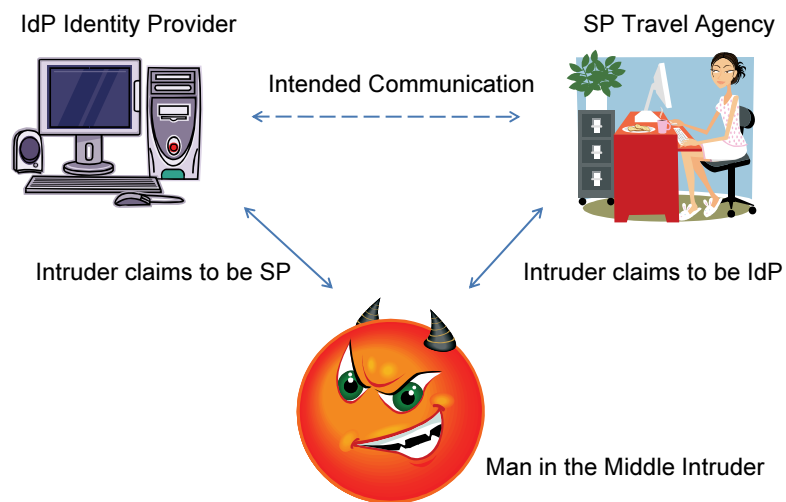
keeping personal data private. This is an improvement over having to repeatedly introduce (and remember) one's credentials. This protocol is related to recent solutions specified by Liberty Alliance, who delegate the task of authenticating a user to an identity provider.

We have provided formal analyses of a protocol that transfers identity and authentication information handled by the systems in an operator's network infrastructure to the applications that implement services provided by third-party service providers. These formal

ios in the process algebra Crypto-CCS. The model checker PaMoChSA (developed by the Security group at IIT-CNR) has then been used to verify their vulnerability with respect to man-in-the-middle attacks by adding a so-called intruder component (see figure).

Asynchronous Web Service Invocation Protocols

Mobile communication networks are typically unstable, since terminal devices may dynamically change their reachability status during their lifetime. Within service-oriented architectures,



Verifying vulnerability by adding an intruder component.

analyses verify the correctness of the protocol and investigate its security properties. The analysis results will be used to identify possible flaws or weaknesses in the protocol and to suggest improvements.

As initial steps towards a full formal security analysis of this protocol, we have modelled a number of user scenar-

ios in the process algebra Crypto-CCS. The model checker PaMoChSA (developed by the Security group at IIT-CNR) has then been used to verify their vulnerability with respect to man-in-the-middle attacks by adding a so-called intruder component (see figure).

towards the successful implementation and evaluation of reliable service-oriented computing applications.

Together with Telecom Italia, we are involved in developing a variant of SOAP coined aSOAP that supports asynchronous communications. This process is driven stepwise by the results of a formal analysis. Hence this activity is different to that described above, where the formal analysis is performed on a protocol already specified in order to verify its correctness. Our aim is to eventually arrive at a proposal of aSOAP of which we can guarantee that it satisfies certain desirable properties. The formal analysis of a preliminary definition of the protocol highlighted several issues and flaws that suggested a complete review was necessary.

We have modelled aSOAP as a set of communicating UML state machines and verified a number of behavioural properties expressed in the action- and state-based temporal logic UCTL. This was done with the model checker UMC, developed in the Formal Methods and Tools group at ISTI-CNR, which creates and traverses the state space on the fly. The advantage of on-the-fly model checking is that often only a fragment of the full state space needs to be generated and analysed to obtain a satisfactory result. The development of UMC is still in progress and a prototypical version is being used internally at ISTI-CNR for academic and experimental purposes. There has not yet been an official public release of the tool, but the current prototype can be accessed via a Web interface.

Links:

SENSORIA:
<http://www.sensoria-isti.eu/>

PaMoChSA:<http://www.iit.cnr.it/staff/fabio.martinelli/pamochsa.htm>

UMC: <http://fmt.isti.cnr.it/umc/>

Please contact:

Maurice ter Beek
ISTI-CNR, Italy
Tel: +39 050 315 3471
E-mail: maurice.terbeek@isti.cnr.it

Challenges in a Service-Oriented World

by Wolfgang Reisig, Karsten Wolf, Jan Bretschneider, Kathrin Kaschner, Niels Lohmann, Peter Massuthe, and Christian Stahl

Interacting services raise a number of new software engineering challenges. To meet these challenges, the behaviour of the involved services must be considered. We present results regarding the behaviour of services in isolation, the interaction of services in choreographies, the exchangeability of a service, and the synthesis of desired partner services.

In a service-oriented world, enterprises use (Web) services to encapsulate parts of their process logic. The Web Service Business Process Execution Language (BPEL) is an established standard to describe Web services. Usually designed in isolation, a service must nevertheless properly interact with other services at run time. A number of techniques have been suggested to check the compatibility of service interfaces and to discover pairs of semantically matching services. The problem of checking behavioural compatibility, however, is rarely addressed. Figure 1 shows typical examples of behavioural incompatibility.

In our joint research groups, and particularly in the Tools4BPEL project, we address behavioural compatibility and provide a series of techniques for detecting and repairing incompatibility, and for synthesizing behaviourally compatible partner services. Here we address the basics of this approach.

Behaviour of Services in Isolation

The designer of a service must guarantee that his/her service is controllable, i.e. there exists at least one properly interacting partner service. Controllability is thus a criterion that is as fundamental as the notion of soundness in the realm of

business processes and workflows. Furthermore, a well-designed service should meet a number of additional properties such as executability of all activities. We suggest techniques for checking this kind of property based on the static analysis of BPEL code and model checking the state space of services.

Interaction of Services in Choreographies

Although a service might be well designed in isolation, it may cause a deadlock in a choreography of services. To prevent such deadlocks, designers must know either how other services in



Figure 1: Services with incompatible behavior (the buyer and seller services wait for each other) and an ill-designed service (a partner of the order service would have to guess the service's capacity).

the choreography behave or how their particular services are supposed to behave. This means public views or operating guidelines can be used. A public view of a service describes communication from the point of view of the service itself. It is an abstracted version of the original service. An operating guideline of a service represents the possible communication behaviours of partners of the service. It is comparable to user instructions shipped with real-world devices. Both public views and operating guidelines support a number of scenarios, in particular service discovery. We propose algorithms for computing public views and operating guidelines.

Exchanging Services

A company may wish to exchange one of its services for another that is better suited for some purposes. The exchanged service should mimic all the interaction capabilities with existing partners. We differentiate between exchange at design time and exchange at run time. While exchange at design time

concerns only new instances of a service, exchange at run time takes care of migrating instances of a service that are already running. As a particular scenario of exchangeability, we study contract-based compositions of services and develop notions of accordance between a contract and some local part of an actual implementation. If all local implementations accord with the contract, the overall composition is guaranteed to behave correctly.

Synthesis of Desired Partner Services

We develop algorithms that synthesize properly interacting communication partners of a given service. These algorithms also support the validation of a given service, test-case generation for a service, and the construction of adapters between services. Adapters can compensate for incompatibilities between services (see Figure 2).

Diagnosis

It would be valuable for a designer to know why a service is not controllable.

Such services typically exhibit static structures that can be conceived as ‘responsible’ for incompatibilities. Knowing the reason helps the designer to remedy the problem. Therefore, we work on diagnosing which structures in a service are responsible for its misbehaviour and on their visualization.

Formal Semantics

In order to be independent of existing service description languages such as BPEL, we suggest a formal Petri-net-based modelling technique for services. The above-mentioned analysis techniques and algorithms are based on our modelling technique. This ensures their long-term applicability and potential use for various service description languages. All algorithms are implemented in our tool Fiona. This tool has a range of capabilities, including generation of operating guidelines and synthesis of communication partners. In addition, service models based on Petri nets are subject to our dedicated model checker LoLA.

We have chosen BPEL as one example of a Web service description language and implemented a compiler BPEL2oWFN, which translates back and forth between BPEL and our service model. Our approach hides the existence of the Petri net model from the users, depicting the analysis results in the BPEL code. Other (Web) service description languages can be analysed by the same techniques, assuming a corresponding compiler.

Links:

<http://www.informatik.hu-berlin.de/top/>

http://www.teo.informatik.uni-rostock.de/ls_tpp/

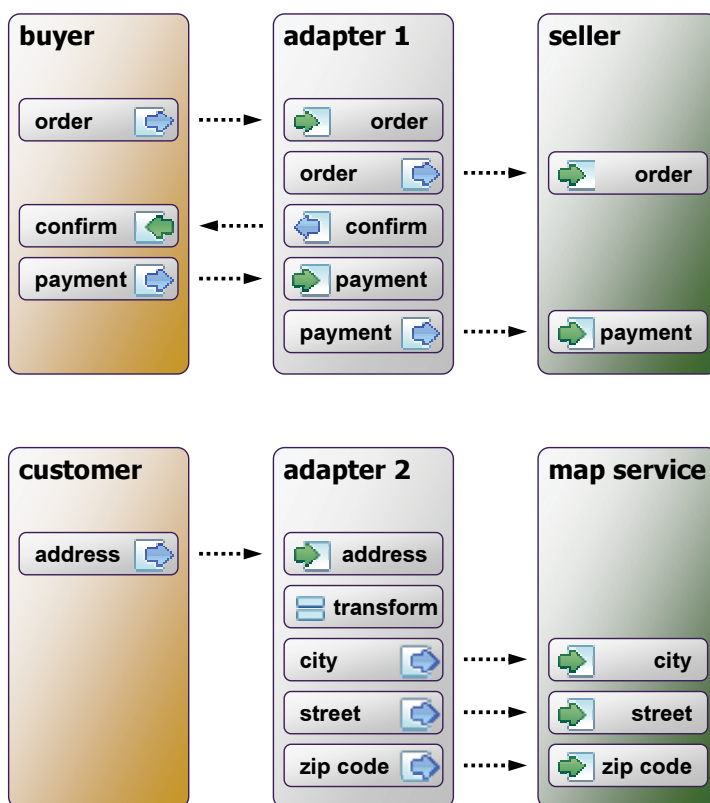
<http://www.informatik.hu-berlin.de/top/tools4bpel/>

Please contact:

Wolfgang Reisig
Humboldt-Universität zu Berlin,
Germany
Tel: +49 30 2093 3066
E-mail: reisig@informatik.hu-berlin.de

Karsten Wolf
Universität Rostock, Germany
Tel: +49 381 498 7670
E-mail:
karsten.wolf@informatik.uni-rostock.de

Figure 2: Adapters can compensate behaviorally incompatible services. Adapter 1 generates a confirmation message to resolve a deadlock. Adapter 2 transforms an address message into its constituents.



RCE and SESIS - Service-Oriented Integration Environment for Collaborative Engineering

by Ottmar Krämer-Fuhrmann

The 'Reconfigurable Computing Environment' (RCE), is a service-oriented software infrastructure for managing collaborative engineering processes. It hides the complexity of heterogeneous and distributed IT systems behind common user interfaces and thus enforces security in the access of data and services. RCE is easily adapted to different application domains: recent work has seen two German ship-building companies build the Ship Design and Simulation System (SE SIS) on top of RCE. RCE was jointly developed by the Fraunhofer Institute for Algorithms and Scientific Computing (SCAI) in Sankt Augustin and the German Aerospace Center DLR in Cologne.

RCE is an open infrastructure framework that offers all the services necessary to operate a distributed collaborative environment. The key features of RCE are its portability and extendability, its security services and that fact that it is application-independent.

RCE is based on OSGi (Open Services Gateway initiative), the industry standard for modular dynamic Java applications. RCE is therefore platform-independent and can be used on any architecture from laptops to mainframes.

RCE can be easily extended by application-specific plug-ins. Services are integrated as plug-ins, the central mechanism known from the Eclipse universe. Non-Java code such as C or Fortran decks can be integrated via wrapper technology, which has already been developed to integrate existing code. This approach allows existing software to be reused and thus earlier investments to be saved.




Figure 1: Distributed RCE architecture.

Communication between the components is realised via a generic interface. There exist several implementations for data transfer, which are deployed

depending on the distance between partners and their security needs. Remote Method Invocation (RMI) realises communication in the same



Figure 2: RoRo ferries by Flensburger Schiffbau Gesellschaft.



Figure 3: Double hull tanker by Lindenau Werft.

compute node, CORBA (Common Object Request Broker Architecture) is used for intranet communication and Web services pass firewalls between company domains.

RCE has an integrated user rights management, which protects access to resources and services. The basic philosophy is that owners control the data or code owned by them, and must grant explicit permission in order for other users or user groups to gain access. This is essential in collaborative industrial environments, where intellectual property rights must be preserved.

The seamless integration of Grid technology into RCE allows the transparent use of resources of service providers. Further, RCE enables engineers to access computational resources on demand, giving them the opportunity to submit computationally intensive simulations to high-performance clusters. Such clusters are usually not available in small and medium-sized companies, because the acquisition of such equipment is not profitable.

In a collaborative network, RCE is deployed as a base system on each computational resource participating in the network. By adding individual plug-ins, each installation can be configured individually: database plug-ins integrate connections to file systems and databases, wrapper plug-ins connect the network to (commercial) simulation codes, and GUI plug-ins allow the engineers to view design data and display results graphically.

SEGIS

Because of its open architecture, RCE can be applied to arbitrary application domains. Cooperation with two German shipyards (Flensburger Schiffbau Gesellschaft and Lindenau GmbH Schiffswerft und Maschinenfabrik) has seen RCE used to integrate consultants and suppliers into the ship design process. The 'Ship Design and Simulation System' (SEGIS) integrates the individual databases involved in the construction of a new ship. The service-oriented architecture allows transparent sharing of data and simulation facilities between the partners. Overall, SEGIS

significantly improves the design process for new ships, which must be constructed in a limited time and involve the expertise of a large number of partners.

DLR and Fraunhofer SCAI are planning to establish RCE as a service-oriented platform for future projects in other application areas. For example, the integration of engineering processes in automotive and aircraft industries is currently in preparation.

Link:

<http://www.scai.fraunhofer.de/kraemer-fuhrmann.html>

Please contact:

Ottmar Krämer-Fuhrmann
Fraunhofer-Institute for Algorithms and Scientific Computing (SCAI),
Germany
Tel.: +49 2241 14 2202
E-mail: ottmar.kraemer-fuhrmann@scai.fraunhofer.de

The ITIA-CNR Customized Shoe Production Service

by Andrea Chiodi, Fabrizio Silva and Andrea Ballarino

The experimental shoe factory in Vigevano allows customers to order their shoes online, with their feet being measured via a 3D digital scan. A Web-based service processes the customer request, accompanied by the foot measurements. The measurements are then matched against existing collections of shoes and a made-to-measure production order is launched.

The Integrated Pilot Plant (IPP) of ITIA CNR is an experimental factory for shoe production designed to handle the whole life cycle of a customizable product. It is located in Vigevano (Italy), a city with a well-known tradition in shoe manufacturing, and is the result of a series of research projects at both Italian and European levels, which have been led by ITIA over the last ten years. These projects focused on the development of Internet-based services, considered strategic for the development of a shoe industry that aims to maintain direct contact with its customers in order to satisfy their specific needs and tastes. The guiding concepts are standardization, collaboration and sustainability.

The production and distribution model is based on a direct relationship between

the customer and the manufacturer: the customer chooses the product, specifies optional features, and supplies morphological information that will enable structural modifications of the end product. Foot measurements are usually taken by means of a 3D foot-scanner based on laser or photogrammetric technology.

The production line is designed to react adaptively to very small (mostly unitary) and heterogeneous production orders. The customer is allowed to check the progress of the order and receives the product at home.

The whole factory acts as a service for on-demand production, operated by Internet-based infrastructure. Orders are accepted either via interactive Web

tools (for direct customer-to-manufacturer purchases), or through client-server negotiation (for purchases operated by a shop assistant using a local retail store software application).

Each shoe model must be described appropriately in order to guarantee optimal fitting and to facilitate mass customization. Information about available products is thus stored in a Web-accessible catalogue of products, which contains an extended set of attributes. These provide not only the usual information such as price, style and size, but also data on measurements and fitting tolerance.

Significant efforts have gone into the definition of standard measurements for both shoes and feet, since conventional

sizing methods provide only approximate measurement information, insufficient for an effective matching. There is a strong need for formal dictionaries of shoe/foot measures, possibly standardized and adopted by the community of shoe manufacturers.

Orders to be produced at the IPP are launched using an Adaptive Production Line: the pair of lasts (the forms on which the shoes are shaped) is first chosen (or built) on the basis of the foot measurements; the components making up the shoe are then modified accordingly; finally the production order is dispatched.

The product catalogue can also include models of shoes proposed by other shoe manufacturers. In this case, the IPP Factory acts as a 'matchmaker' Web service between the customer's requirements (choice of shoe and size of feet) and the products offered by different manufacturers. The manufacturer who produces the relevant shoe model is able to take advantage of the technical information collected by the IPP service during the selection and ordering of the shoes.

The whole system makes extensive use of service-oriented technology like



The IPP Factory - production line and control deck.

SOAP and WSDL. The full production cycle is maintained through an internal Web-based infrastructure, supporting various interactive and client-server Intranet services, and enabling cooperation between the several modules of the highly automated adaptive production line.

Some of these Web services are also made accessible to suppliers and other external collaborators, to allow easy exchange of information during produc-

tion. Efforts are under way to enhance these services by adding traceability features to the products, thereby covering the full life cycle from cooperative design to delivery, use and dismissal.

Link:

<http://www.itia.cnr.it/>

Please contact:

Andrea Chiodi
ITIA-CNR, Italy
E-mail: a.chiodi@itia.cnr.it

The ToolBus: A Service-Oriented Architecture for Language-Processing Tools

by Paul Klint

The paradigm of service orientation creates new opportunities for language-processing tools and Interactive Development Environments. At CWI we have developed the ToolBus, a service-oriented architecture with application areas like software renovation and implementation of domain-specific languages.

Service-oriented architectures aim to decouple processing steps. They are usually applied in the context of heavy-weight business applications where traceability and transaction management are the dominant requirements.

In language-processing tool suites and Interactive Development Environments, other requirements prevail: for instance, the tight integration of tools and efficient exchange of data are more prominent as requirements. At CWI we have developed the ToolBus, a service-oriented architecture, to achieve this. It is based on process algebra as a concur-

rency paradigm and on ATerms as a data exchange mechanism. The ToolBus forms the foundation for a suite of language-processing tools combined in the ASF+SDF Meta-Environment. ASF+SDF is a term rewriting language that extends the syntax definition formalism (SDF).

There are many urgent language-processing tasks that require a quicker answer than can be achieved by building a dedicated tool from scratch. Examples are performing a domain-specific analysis ("Is the memory management API used consistently?"), executing a dedi-

cated transformation ("Refactor this code to use the new API.") or building support tools for a new domain-specific language. In all these cases, the best approach is to combine existing tools with newly written ones in order to solve the problem quickly.

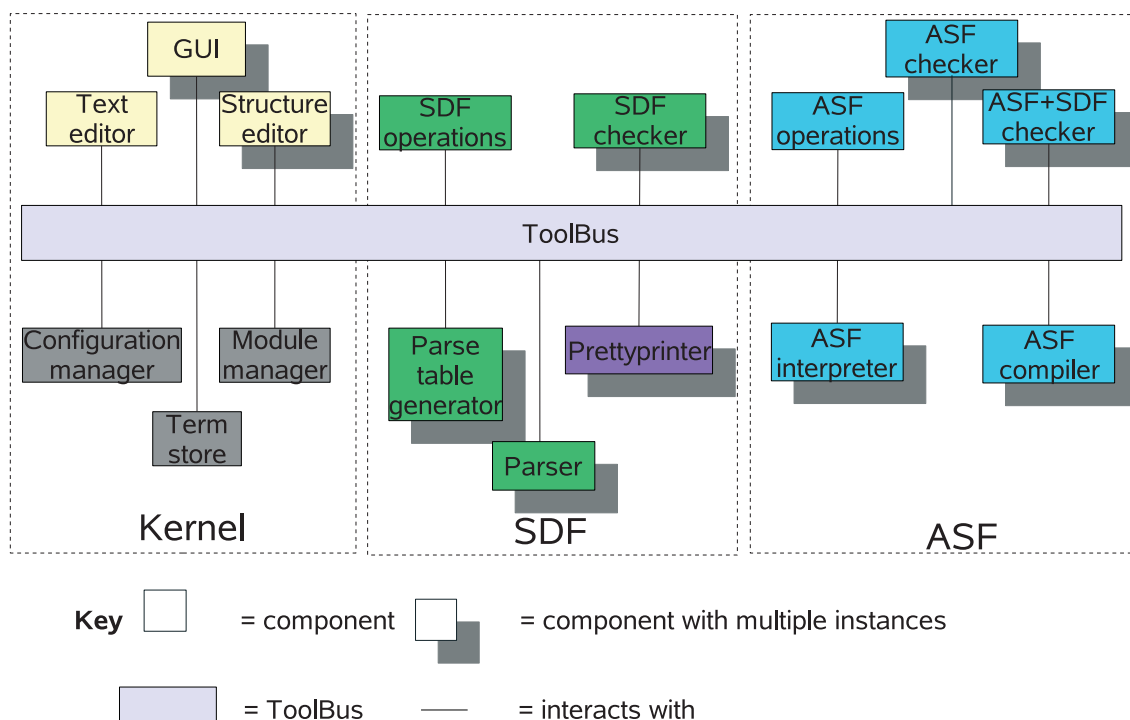
The tools typically available in the language engineer's tool chest are generators (for parsing, formatting and code generation) and generic tools for editing, user-interfacing and visualization. These tools are usually written in different languages and run on different platforms.

The relevance of service orientation to this domain is evident but we apply it with some twists. Based on experience, we have come to the conclusion that the straightforward use of XML for exchanging intermediate data is inadequate: the parse trees that arise from analysing hundreds of thousands of lines of code simply become too bulky. In order to solve this problem we represent intermediate data as ATerm (short for Annotated Term), a directed acyclic graph that maximizes subterm sharing and can be represented very concisely.

the Meta-Environment. Observe that variations in the Tscript lead to variations in the resulting system; product families can thus easily be supported.

Recently the ToolBus entered a new phase in its development. The existing ToolBus was implemented in C; we are now about to finish a reimplementation in Java in order to profit from Java's better structuring facilities and from the direct availability of many relevant communication and (Web) service libraries. This new version will also

various academic and industrial parties for software analysis, software transformation and domain-specific language development. Examples are analysis and refactoring of the C code for ASML's lithography machines, the renovation of administrative Cobol code by Getronics, and the use of a financial domain-specific language by the Fortis bank. Current work on the ToolBus is done in cooperation with Technical University Eindhoven and University of Amsterdam. Within ERCIM, we cooperate with INRIA.



The ToolBus while orchestrating tools in the Meta-Environment.

By providing all the relevant tools with an ATerm interface, huge amounts of data can be shipped between tools while sharing is preserved.

Another issue is how to orchestrate the execution of all these tools. To this end, we connect them to the ToolBus, which can simply be described as a programmable, ATerm-enabled, service bus. The orchestration is described by Tscript, a scripting language based on process algebra that supports parallelism, asynchronous and synchronous communication and tool control. This allows the construction of large, heterogeneous and distributed applications. The figure illustrates the use of the ToolBus while orchestrating the tools in

address issues such as built-in profiling and monitoring, efficient tool-to-tool communication, and better isolation and recovery of malfunctioning tools. In addition to its existing service and networking capabilities, this new implementation will also enable execution on multi-core computers. In this way, the whole distribution spectrum - from wide area to multi-processors on a chip - can be handled. This enables applications in which high-density local computation clusters are loosely coupled via a wide-area network. We envisage that more advanced software analysis tasks will require such an infrastructure.

The ToolBus is distributed as part of the Meta-Environment, which is in use by

Links:

<http://www.meta-environment.org/>
<http://www.cwi.nl/~paulk/>

Please contact:

Paul Klint, CWI, The Netherlands
 Tel: +31 20 592 4126
 E-mail: Paul.Klint@cwi.nl

Open Grid Services for Improving Medical Knowledge Discovery

by Manolis Tsiknakis

ACGT (Advancing Clinico-Genomic Trials on cancer) is an FP6 integrated project focusing on the development of a semantic Grid infrastructure to support multicentric, post-genomic clinical trials. This will enable discoveries in the laboratory to be quickly transferred to clinical management and the treatment of patients.

Recent advances in research methods and technology have resulted in an explosion of information and knowledge about cancers and their treatment. Exciting new research on the molecular mechanisms that control cell growth and differentiation has resulted in a significant improvement in our understanding of the fundamental nature of cancer cells, and has suggested valuable new approaches to cancer diagnosis and treatment. Despite these advances, the lack of a common infrastructure has prevented clinical research institutions from mining and analysing disparate data sources. This inability to share technology and data developed by different organizations is severely hampering the research process.

Post-Genomic Clinical Trials

The ACGT project has been structured to deal with this problem. The ERCIM office is the administrative project coordinator, while FORTH is responsible for scientific coordination. The project has selected two cancer domains (breast cancer and Wilm's tumour or pediatric nephroblastoma) and has defined specific trials which are feeding the requirement analysis and elicitation phase of the project. A third trial is also included, which focuses on the reuse of multilevel biomedical data produced in the previous two trials and the integration of advanced technology (including interactive visualization, virtual reality technology and in silico tumour growth simulations). Here the objective is to explore simulated predictions of tumour growth and treatment response.

The Breast Cancer Trial

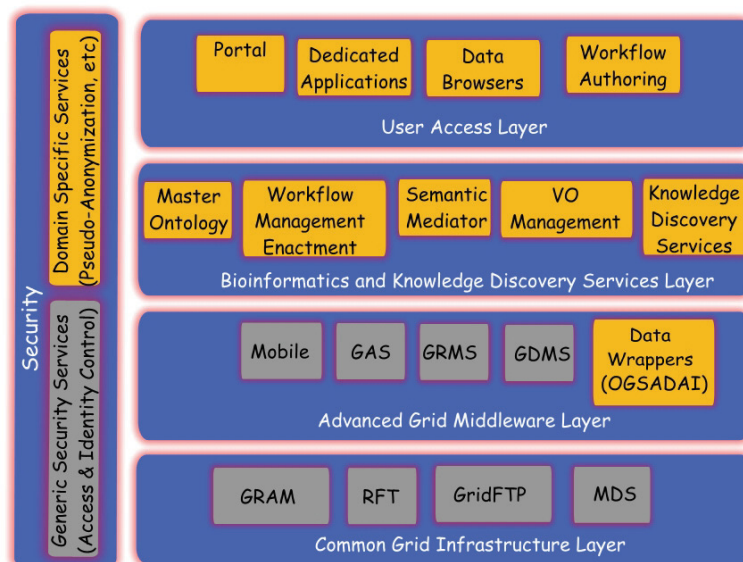
Breast cancer is both genetically and histopathologically heterogeneous, and the mechanisms underlying its development remain largely unknown. The ACGT Test of Principle (TOP) study aims to identify biological markers associated with pathological complete response to anthracycline therapy (epirubicin), one of the most active

drugs used in breast cancer treatment. Supported by in vitro and preliminary in vivo data, this study is designed to test prospectively the value of topo II alpha gene amplification and protein overexpression in predicting the efficacy of anthracyclines.

These clinical trials are multicentric (many different research organizations

to be embedded in the platform, and to relevant ethical issues, thus creating optimal conditions for service uptake.

Since we see the requirements engineering process as a structured set of activities that will lead to the fulfillment of the final system requirements, an iterative method has been adopted, based mainly on scenarios and prototyping.



The ACGT layered functional architecture.

are participating), and post-genomic, meaning they require generation, management, integrated access, processing and analysis of multilevel biomedical data, including transcriptomic, proteomic and imaging data.

As a result, the ultimate objective of the ACGT project is the development of semantic Grid infrastructure that offers high-level tools and techniques for the distributed mining and extraction of knowledge from data repositories available on the Grid. This infrastructure will make use of semantic descriptions of components and data and will offer knowledge discovery services in the domain of cancer research. Special emphasis is given to the trust that needs

Explicit scenarios have been developed that represent documented user needs and also provide a technology-driven description of the requirements of the system under design, as understood by experienced technological experts.

Initial System Architecture

From a detailed analysis of documented user requirements, it is apparent that a complex technical infrastructure must be developed if support for integrated access, analysis and visualization of multilevel, heterogeneous data is to be provided. A detailed analysis of the scientific and functional requirements of the ACGT infrastructure was performed, together with an analysis of the current

state of the art in terms of technological infrastructure, data resources, data representation, exchange standards and ontologies.

With respect to the state of the art, the myGrid project (<http://www.mygrid.org.uk>) is focusing on providing in silico support for experimental research, while the cancer Biomedical Informatics Grid (caBIG - <https://cabig.nci.nih.gov/>) is creating a virtual community within which resources can be shared and the key issues of cyber infrastructure tackled.

From a technical point of view, the requirements identified can be met using a federated, multilayer, service-oriented and ontology-driven architecture. The ACGT project decided to build on open software frameworks based on the WS-Resource Framework (WSRF) and Open Grid Service Architecture (OGSA), which are the de facto standards in Grid computing. These standards are implemented in the selected middleware, namely Globus Toolkit 4 (GT4) (<http://www.globus.org>) and Gridge (<http://fury.man.poznan.pl/gridge/>).

An overview of the ACGT system layered architecture is given in the figure, and includes the following layers:

- *Common Grid Infrastructure Layer:* this comprises the basic 'Grid engine' for accessing remote resources in a Grid environment. It provides a com-

mon interface for Grid resources used by higher-level services

- *Advanced Grid Middleware Layer:* this comprises advanced Grid services, which operate on sets of lower-level services to provide more advanced functionality
- *Bioinformatics and Knowledge Discovery Service Layer:* this includes all the ACGT-specific services, such as the ACGT Master Ontology, the Clinical Trial on Cancer Metadata Services, semantic mediation services and distributed and privacy-preserving data-mining and knowledge discovery services
- *User Access Layer:* this allows users to realize complex biomedical applications by combining basic services from the underlying layers and exploiting the resources and data provided by the research centres that form different CT Virtual Organizations (VOs)
- *Security Layer:* access rights, security and trust-building are issues addressed by this layer.

Biomedical Grid Intelligence

In a 'Grid-enabled' data-sharing VO, datasets may not be well known amongst all VO participants. To integrate highly fragmented and isolated data sources, we need semantics in order to answer higher-level questions. It therefore becomes critically important to describe the context in which the data was captured. We describe this

contextualization of the data as metadata. Semantic integration in ACGT thus relies on metadata publishing and ontologies.

We see as our main future research challenge in ACGT the development of an infrastructure that is able to produce, use and deploy knowledge as a basic element of advanced applications. This will mainly constitute a Biomedical Knowledge Grid. Metadata is critical to achieving such an objective. We use OWL-S to develop metadata and service ontologies for describing Grid Services so that they might be discovered, explained, composed and executed automatically.

Our initial investigations have also revealed the need for a sophisticated model of provenance, since the use of both elementary and advanced workflows (workflows containing other workflows), is becoming a very important goal in our R&D work. This requirement also involves maintaining complex metadata relating to workflows in the ACGT Grid middleware.

Link:

<http://www.eu-acgt.org>

Please contact:

Manolis Tsiknakis

ICS-FORTH, Greece

Tel: +30 2810 391690

E-mail: tsiknaki@ics.forth.gr

Application of SensorWebs for Intelligent Satellite Tasking

by Adrian Grenham

The ultimate vision for the Sensor Web is to use service-oriented architecture (SOA) techniques and the ubiquitous power and utility of the Web to integrate billions of static and mobile sensors that are already present in the environment. The resulting capability will release an unprecedented wealth of environmental monitoring information at all levels: global, national, regional and local. In order to investigate and demonstrate one aspect of this vision, SciSys research is working on the automatic triggering of satellite-based observations by in situ alerts.

The value of earth observation satellite imagery in monitoring changes in the environment is largely dependent on the timeliness of the data acquisition. Models routinely used to forecast weather and related natural phenomena such as floods, employ sophisticated numerical analysis techniques to provide forecasts from hours to many days

in advance. Many of these models rely on information from large numbers of (often heterogeneous) in situ sensors; these collect diverse data including soil moisture, river levels and peak flow, coastal tide conditions and rainfall. Access to in situ sensors tends to be ad hoc, employing widely different protocols and data formats.

The goal of the Sensor Web is to provide a network that links ground and space-based instruments for autonomous and collaborative observation services. Ideally this is achieved by the seamless integration of software and communications links between sensors. Typically models raise alerts concerning events of interest; an observation campaign is




Figure 1: Using multi-temporal SAR images, the vast expanses of sandbank covered by the sea at high tide can be mapped, clearly showing the orientation of the river channels (reproduced with permission. ©LCC and BNSC).

then triggered that employs an ad hoc and dynamic set of sensors. Subsequently, data is acquired as quickly and in as much depth as possible in a given time period. In this example a satellite instrument can be tasked to provide imagery for constructing before and after sequences to aid situation awareness, defence planning and relief activities if needed.

The aim of the Open Geospatial Consortium (OGC) Sensor Web Enablement (SWE) specifications is to create Web-based sensor networks that will make all sensors and sensor data repositories discoverable, accessible and (where applicable) controllable via the worldwide

Web. It is a revolutionary approach to exploiting Web-connected sensors, and represents an extension of the concepts widely demonstrated by the Web Map Service (WMS), Web Feature Service (WFS) and Web Coverage (WCS) specifications that facilitate interoperability in geospatial Web services.

Besides the principles governing Web services, the OGC specifications describe the functionality of the Sensor Web:

- discovery of sensor systems, observations, and observation processes that meet the immediate needs of an application or user
- determination of a sensor's capabilities and quality of measurements

- access to sensor parameters that automatically allow software to process and geo-locate observations
- retrieval of real-time or time-series observations and coverage in standard encodings
- tasking of sensors to acquire observations of interest
- subscription to and publishing of alerts to be issued by sensors or sensor services based upon certain criteria.

The specifications allow Web-resident sensors to be discoverable and accessed through a set of self-describing software services. A dedicated profile of XML, SensorML provides a generic description of any sensor type that can be queried for any accessible sensor. The profile includes modelling sensor response as a process that may be described as a data transform specific to a physical device (eg an actuator, detector or filter). Individual processes can be combined into more complex processing chains that satisfy some real-world problem.

SciSys is involved in a European Space Agency project to test intelligent satellite tasking. The satellites are triggered by alerts raised by a Sensor Web based on a prototype implementation of the OGC specifications. The project comprises an international consortium including GIM (Belgium), VITO (Netherlands), SpaceBel (Belgium), Soldata (Spain) and GeoID.

The project will deliver a flood-monitoring service integrated into the ESA-ESRIN Service Support Environment

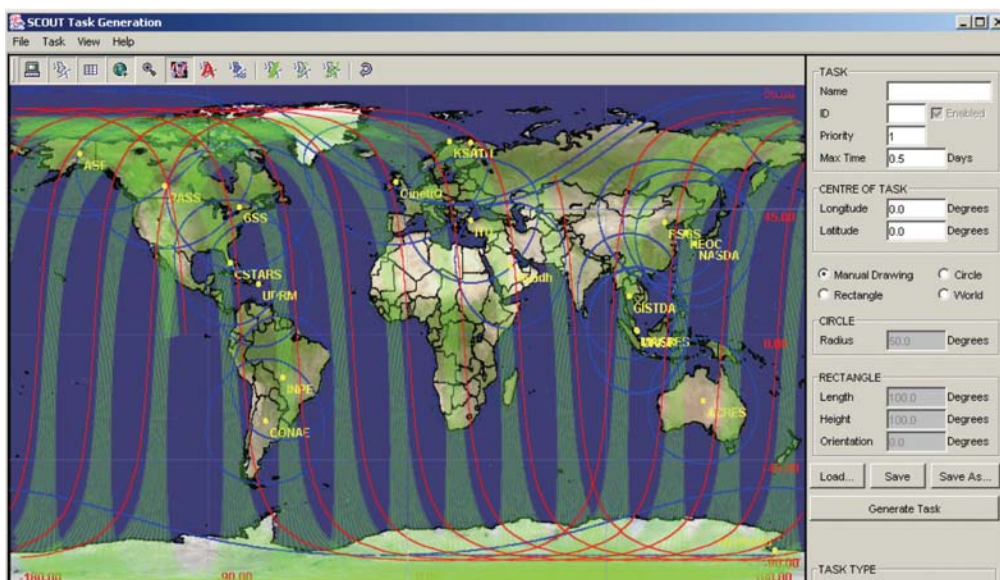


Figure 2: Determining imaging opportunities using the ground coverage of an LEO satellite.

(SSE). Users will be able to subscribe to receive flood alerts issued for a defined geographical area of interest. The service will provide users with a set of satellite-imaging opportunities by querying an SWE Sensor Planning Service (SPS) that acts as the asset controller for a number of satellite SAR missions. The orbit description, sensor and tasking details for each mission are used to determine the earliest opportunity to image the target area. The SPS 'submit' operation formats a tasking request including all the parameters required by the relevant mission-planning system.

The prototype testbed will be used to simulate the provision of timely SAR images of Morcambe Bay in northwest England. These can show the orientation of highly mobile river channels prior to forecasted flood events, adverse weather and seasonal high tides. Mapping the channels in advance of exceptional tide conditions makes it possible to assess the location and effectiveness of sea defences and to alert populations likely to be affected.

Currently such satellite tasking remains a largely manual process and requires

fairly long lead times. However, with warnings issued 6 to 48 hours prior to an event there is an obvious potential to acquire timely data through improved responsiveness to flood alerts.

Links:

SciSys: <http://www.scisys.co.uk>

OGC: <http://www.opengeospatial.org>

Please contact:

Chris Lee

Head of Sales and Marketing

(Space Systems), SciSys Ltd, UK

E-mail: chris.lee@scisys.co.uk

Open-Source SOA for Small and Medium Organisations

by Pascal Bauler and Nicolas Biri

Recent and planned activities of the Centre de Recherche Public – Gabriel Lippmann include the design of open-source service-oriented architectures (SOA) and the implementation of these new concepts in small and medium organisations.

Most modern organisations rely on complex IT environments that play a critical role in their daily business. Most of these systems have been designed over the last few decades and rely on heterogeneous technology, including mainframe applications, client/server solutions, multi-layer applications, off-the-shelf software and custom developments.

In general, the subsystems are stand-alone solutions and do not present an integrated IT environment. As a consequence, data synchronisation between the various subsystems usually requires manual interventions and is prone to error.

The project goal is to work out a generic approach to modernising the IT environment of Small and Medium Organisations (SMOs), by migrating to a service-oriented architecture (SOA). The proposed solution relies on open standards, which make the various sub-components exchangeable and facilitate the exploitation of the SOA.

As the IT budgets of SMOs are usually very limited, the proposed SOA is based on multiple open-source projects, thus reducing licensing costs. The proposed SOA solution has a three-layer architec-

ture. The technical integration layer consists of various connectors establishing a technical integration between the SOA and the external applications. An Enterprise Service Bus (ESB) makes the various subsystems independent of transport protocol by hosting data transformation and message routing services. Finally, the orchestration layer provides the SOA with a user interface, offering business process execution facilities that seamlessly access the underlying heterogeneous applications.

The success of an SOA project strongly depends on the ease of use of the newly proposed architecture. Consequently, the SOA should facilitate the addition and publishing of new services as well as the exploitation of existing services. In order to achieve this goal, model-driven software design techniques are used to generate code segments that accessing existing services, combined with a dedicated framework used to design new SOA services.

Technical Considerations

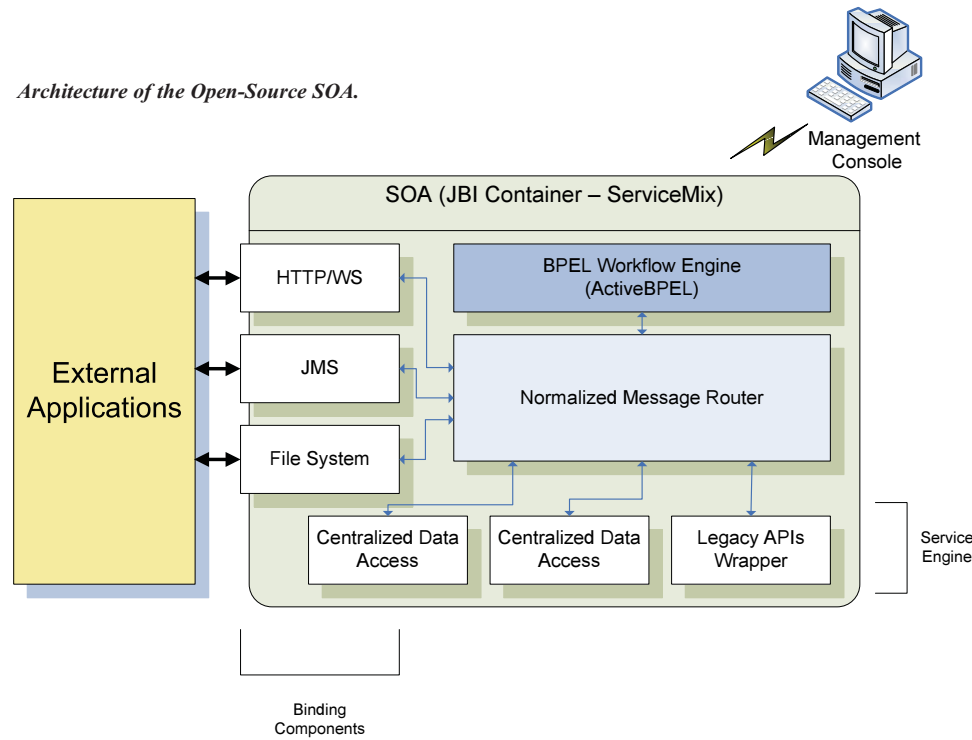
As mentioned above, the main project goal consists in working out a practical approach to modernising the IT environments of SMOs by migrating to a standards-based service-oriented architecture that relies on open-source technol-

ogy. In order to achieve this goal and after comparing several solutions, ServiceMix, an open-source service bus based on Java Business Integration (JBI) was selected and systematically extended (see Figure 1). Newly designed JBI components guarantee protocol independence and the orchestration layer is realized by the open-source version of ActiveBPEL, which was adapted to act as a native JBI component. This technology represents the technical back-end of the SOA solution and is enhanced by a Java-based component framework. This is accessible through a dedicated Eclipse plug-in used by the development teams while creating and deploying new JBI components or while integrating additional applications with the SOA. Access to existing services managed by the SOA is provided by a newly developed domain-specific language (DSL), which is used by the OpenArchitectureWare Model Driven Systems Development (MDS) framework to generate the appropriate application code.

Practical Use Cases and Perspectives

Integration projects must prove their worth in practice. The proposed solution is currently being deployed in a research project running in collaboration with the

Architecture of the Open-Source SOA.



Luxembourg National Family Benefits Fund (CNPF). Initial results involving the integration of mainframe applications, WinDev proprietary software and Java subsystems are very promising. A first application relying on the new SOA will automatically handle the family allowances for French border commuters and is scheduled to be released in May 2007.

In addition to the step-by-step migration of the Luxembourg National Family Benefits Fund to the SOA, the same platform will be used in collaboration with the Luxembourg Chamber of

Commerce, but will be extended by enhancing its security aspects and by adding service repository functionalities. At the same time, the potential business benefits resulting from scientific exploitation of the heterogeneous data available throughout the new architecture identify the SOA project as a good starting point for subsequent Business Intelligence (BI) projects.

Conclusion

The proposed solution integrates and enhances several open-source projects and highlights the easy exploitation of the SOA. First practical results have

been obtained in collaboration with the CNPF. The mid-term goal is to further extend the SOA by working on its security aspects, service repositories and business intelligence systems. In parallel, the SOA will be used in the context of other research projects.

Please contact:

Pascal Bauler and Nicolas Biri
 Centre de Recherche Public –
 Gabriel Lippmann, Luxembourg
 Tel: +352 47 02 61 1
 E-mail: bauler@lippmann.lu,
 biri@lippmann.lu

FUSION: Business Process Fusion Based on Semantically Enabled Service-Oriented Business Applications

by Spiros Alexakis and Alexa Schumacher

Small and medium enterprises (SMEs) cooperating with international partners in the enlarged Europe need holistic Enterprise Applications Integration (EAI) solutions in order to operate effectively. At the same time, they face intercultural barriers, since current interoperability and integration efforts are focused on data rather than on processes.

The aim of the FUSION project is to support collaboration and interconnection between commercial enterprises by developing a framework and innovative technology that allows the semantic fusion of heterogeneous service-ori-

ented business applications. In particular, FUSION has a three-fold focus:

1. The development of an innovative approach, methodology and integration mechanism for the semantic inte-

gration of a heterogeneous set of business applications, platforms and languages within SMEs.

2. The integration of research activities carried out in the enlarged Europe in

the areas of Business Process Management, Semantic Web and Web services.

3. The validation of research results by developing proof-of-concept pilots in collaborative commerce growth across semantically enriched value networks across the enlarged Europe.

FUSION will facilitate three transnational business cases, typical examples of cross-organizational collaboration. The first is the integration of transactions of a franchising firm (Greece, Poland, Romania, Bulgaria, Ukraine, Cyprus and FYROM); the second pilot deals with the automation of international career and human resource management services (Hungary and Germany); the final example is a collaboration of companies in a chain of schools of foreign languages and computing (Bulgaria, FYROM, Albania).

The research project FUSION is led by SAP AG (Germany). Research activities are coordinated by the Institute of Communication and Computer Systems (ICCS, Greece). Pilot case execution is led by CAS Software AG (Germany), while validation of the research results is coordinated by the South-East European Research Centre (SEERC, Greece). In total, the FUSION consortium consists of fifteen partners from five European countries (Germany, Poland, Greece, Hungary and Bulgaria), including research institutes,

technology developers, innovation transfer bodies and end users.

The FUSION solution will invoke the creation, administration and deployment of Web Services Instances of the pre-selected features of the enterprise applications and their semantic description (Semantic Profile). This is based on a business concept model called FUSION ontology that serves as a common reference and allows the semantic integration of the business applications. The deployed Web Services Instances will be published in a semantically enriched UDDI service registry, while the semantic profiles created for Web Services Instances will be registered in the FUSION semantically enriched registry. The latter is part of the FUSION Semantic Repository, powering the system-integrated categorization and discovery services.

We envision that in working with the FUSION solution, business analysts will:

- define the concepts that exist inside the business application (eg product, contact, order etc) and use well-defined business ontologies/concept models that are independent from the technical architecture of the business application
- associate the concepts and services with the business application repository/resources
- create and administrate Web Services Instances, create semantic profiles of Web Services Instances and publish

the profiles on the FUSION semantically-enriched registry

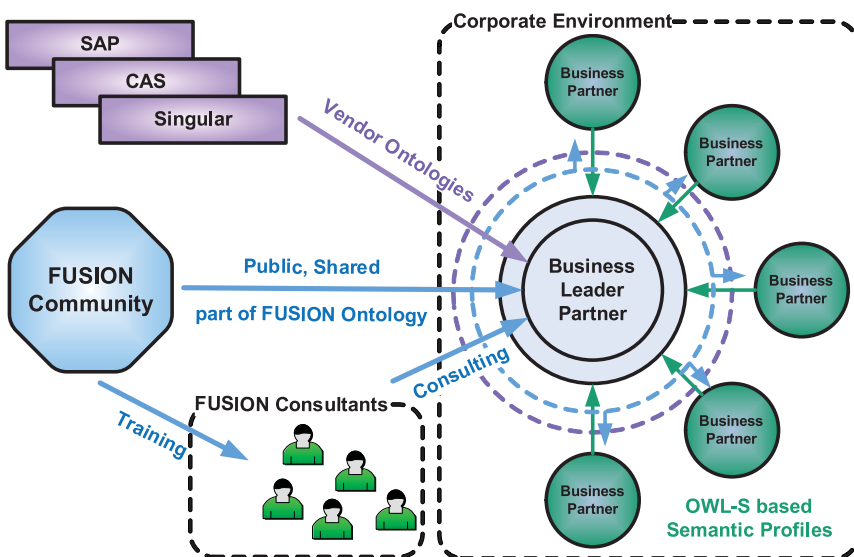
- orchestrate aggregate compositions of Semantic Web services based on FUSION-enabled descriptions of various business applications.

To this extent, those involved in the FUSION System life cycle are IT consultants (who are responsible for extending the ontology with necessary concepts and annotating the Web services), business process consultants (who create generic processes that can be customized for more specific installations), and enterprise application software vendors (who adapt their products to the FUSION approach).

It is expected that the project results will include the following:

- the FUSION approach for semantic service-oriented Business Application integration covering essential business processes between collaborative organizations
- the FUSION methodology for semantic service-oriented Business Application Integration that will facilitate the integration of business software applications
- the FUSION integration mechanism will simplify the interconnection of heterogeneous information systems, resource sharing and services provision
- three FUSION cases will prove the concepts and tools of the solution, concerning three different enlarged Europe scenarios.

The project is supported by the European Commission (IST-027385).



System Development Scenario.

Links:

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

FUSION ontology:

<http://www.imu.iccs.gr/software/enio>

Please contact:

Spiros Alexakis

CAS Software AG, Germany

E-mail: Spiros.Alexakis@cas.de

VantagePoint: Seeing is Believing

by Jarmo Kalaoja, Julia Kantorovitch, Toni Piirainen and Ilkka Niskanen

Service-oriented architectures, Web services, Semantic Web, ontologies, OWL, RFD, OWL-S... any chance for the service developer and original user (eg researcher) to survive in understanding of the complex world of today's service-related semantic technology? Researchers in the VTT Software Architectures and Platforms department are developing a VantagePoint tool, which will help designers and developers to better understand semantic service-oriented architectures (SOAs). An understanding of semantics will enable application designers to choose appropriate ontologies for their applications, and to determine which semantic-based approaches are beneficial and what can be easily integrated with semantic reference architectures.

Pervasive computing is the trend towards increasingly ubiquitous, connected computing devices in the environment, and is being brought about by a convergence of advanced electronic, wireless technology with the Internet. Pervasive computing devices can be mobile or embedded in almost any type of physical environment (eg home, office or car) imaginable, and all communicate through increasingly interconnected networks.

It is expected that in the future, 'smart' devices all around us will maintain current information about their locations and the contexts in which they are being used, as well as relevant data about the users. Computing will become so naturalized within the environment that people will not even realize that they are using computers.

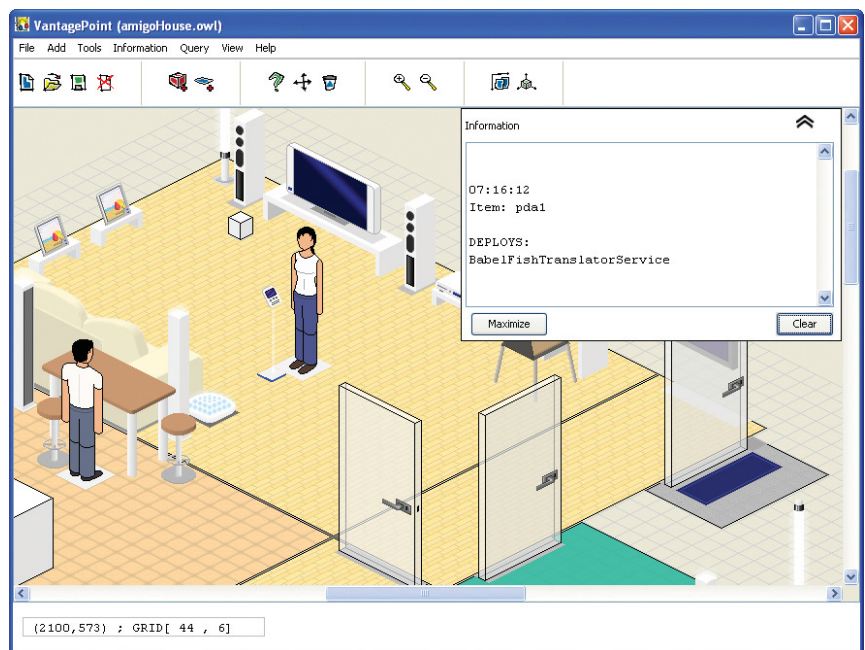
To this end, researchers are looking for innovative concepts and approaches. Among the emerging technology expected to prevail in the pervasive computing environment of the future are service-oriented architectures (SOAs) and semantic ontology-based technologies. SOAs provide a means of resolving the complexity, heterogeneity and dynamicity issues, abstracting the complexity of involved systems and enabling loose and dynamic coupling among heterogeneous resources through automated service publication and runtime service discovery and composition. Semantic ontology-based technologies (eg RDF, RDF-S, OWL, OWL-S, ontologies and reasoning tools) promise far more effective machine-to-machine communication. The use of ontologies enables computational entities and services to use a common set of concepts and vocabularies to represent knowledge about a domain of interest, and to interact with each other. This allows the relationships between entities to be

more clearly expressed and thus better reasoning on their properties.

In the IST Amigo project, which aims to develop an open networked home system as part of the ambient intelligence pervasive computing vision, we adapt semantic modelling approaches to

cutted and the preferences specified in user profiles, may assist in dynamically selecting the services that best meet user needs.

However, before such reference service discovery architecture can be implemented, effectively used and widely



The isometric view of Amigo House.

effectively integrate, discover and compose heterogeneous services and devices in a ubiquitous and seamless manner in order to address user requirements. Web Services and OWL-S are adopted as a convenient paradigm for service representation in Amigo. The purpose of ontologies is to enrich service descriptions, including both the functional (service capabilities, inputs, outputs) and non-functional properties of services. The non-functional properties of a service, such as quality of service (QoS), the context in which it is ex-

accepted, the ontology-based semantics for networked devices and services must be well understood. An understanding of semantics will enable application designers to choose which ontologies can be used in their applications. An understanding of how semantics can be used will help application and service developers to determine which semantic-based approaches will be beneficial to the support of their applications, and what is required to be easily integrated in such reference architecture.

To offer a comprehensive response to the above concerns, researchers in the VTT Software Architectures and Platforms department are developing a VantagePoint tool, which will make the contextual semantic information related to service descriptions easier to understand, and its use by an application developer foolproof. The tool allows users to view ontology instances associated with complex contextual information in a more illustrative and comprehensible way. It also allows users to semantically model and interactively simulate contextual environments of interest. These may be either physical real-world (ie devices, services, functional capabilities of service, contexts) or conceptual (business boundaries, networking or security domains). It supports the conceptual design of appli-

cations (eg verifying a service composition logic) or middleware-level services (eg semantic service discovery) against one or more contextual scenarios. Moreover, we believe that the research presented is a step toward the better understanding and wider acceptance of ontology-based semantic technology also for non-Web services.

VantagePoint is written in Java; it uses the Jena interface to manage OWL ontologies, and Java 2D graphics to visualize them. VantagePoint can have several visualization libraries containing domain-specific icons. These libraries are stored as simple text files that contain URLs of the icon files providing isometric visualization from different perspectives (PNG images), and a URL of a semantic class description in

one of VantagePoint's semantic libraries. A browser tool is provided for examining the visualization libraries. While the current libraries relate to intelligent home applications, future elements will describe other intelligent environments such as car, plane, or mobile outdoor domain.

Links:

<http://www.hitech-projects.com/euprojects/amigo/>

<http://www.vtt.fi/proj/vantagepoint/>

Please contact:

Julia Kantorovitch, Jarmo Kalaoja
VTT, Finland

Tel: +358 20 722 2334

E-mail: Julia.Kantorovitch@vtt.fi,
Jarmo.Kalaoja@vtt.fi

Diapason: An Engineering Environment for Designing, Implementing and Evolving Service Orchestrations

by Frédéric Pourraz and Hervé Verjus

The aim of Diapason is to allow designers of service-oriented architectures (SOAs) to precisely orchestrate services using an orchestration language based on pi calculus. Once defined, an orchestration can be verified against properties and constraints, can be deployed as a new service and can evolve dynamically and on the fly.

SOAs are service-based applications for which classical software engineering approaches fail. This is due to the heterogeneous, autonomous, widely distributed and loosely coupled nature of the services. In other words, when building an SOA, the designer does not control the service's implementations. As SOAs are increasingly used to support widely distributed software-intensive systems in a plethora of domains (business, manufacturing, health, Grid-based applications, military etc), the design, implementation and evolution of SOAs is a challenging problem.

Diapason Approach

We introduce a new environment called Diapason. Diapason is an SOA-based systems-engineering environment that supports service orchestration, along with the analysis, deployment, execution and evolution of that orchestration. Diapason provides a layered formal language called pi-Diapason, which relates

to SOA structure/topology with most of the service orchestration patterns already proposed (www.workflowpatterns.com). It is also extensible (ie architects can define their own extensions and concepts) and executable.

It also provides:

- a properties definition language called logic-Diapason
- an analyser (a properties checking tool)
- an animator (a graphical simulation tool)
- deployment mechanisms
- a virtual machine (a runtime engine).

Basically, services provide operations that can be remotely invoked. The manner in which the services (and operations) are implemented is unimportant, since services are considered to be black boxes that can be composed (orchestrated). When orchestrating Web services, WSDL files are referenced in order

to obtain the signatures of the operations and to then invoke these operations at runtime.

Once defined, a service orchestration can be deployed as a new service that can be further reused in other orchestrations and/or can be modified according to new requirements. When an orchestration is reused it is considered to be a standalone service, providing its own operation(s). Hence, the orchestration can be composed with other services and be involved in other orchestration(s). A service orchestration is executed by way of a virtual machine that interprets pi-Diapason language. In this way, pi-Diapason can be used as a formal and executable SOA specifications language.

Service composition is a key issue in Diapason. Pi-Diapason is formally based on polyadic high-order Milner's pi-calculus, which is a form of process algebra. Service orchestration descrip-

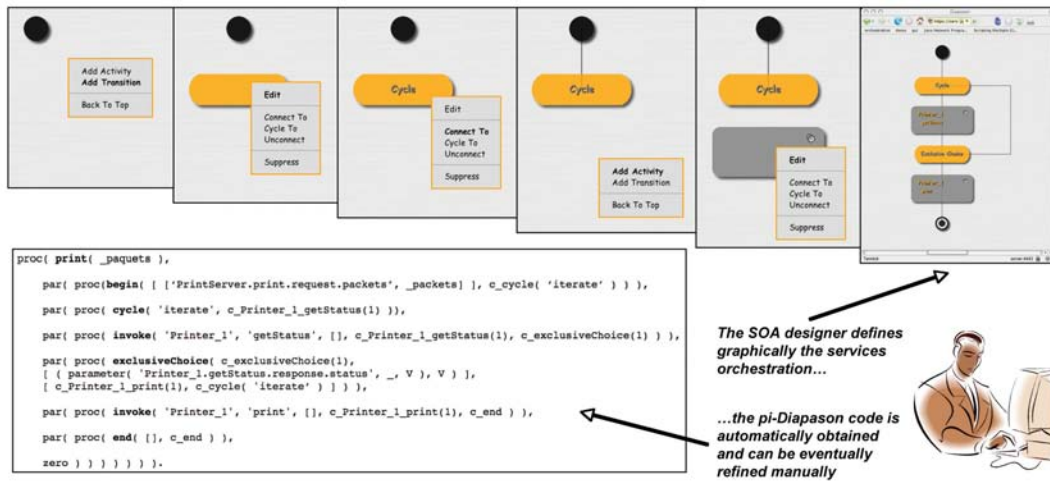


Figure 1: Orchestration description (Diapason graphical editor).

tions use orchestration patterns that are formally defined in terms of pi-calculus processes. Basically, a service orchestration is a complex pi-calculus process that composes other pi-calculus processes. In order to simplify the description of an orchestration, Diapason provides a lightweight and intuitive graphical editor.

Thanks to the pi-calculus mobility (introduced in the first-order pi-calculus but extended to behaviour mobility support in the high-order pi-calculus), we can dynamically modify the service orchestration at runtime without stopping the execution of this orchestration. Evolving a service orchestration at runtime is also very challenging. Evolving a Diapason service orchestration is the same as evolving a pi-calculus process.

Pi-Diapason provides pi-calculus-specific constructs that are responsible for evolution: using these constructs, the SOA architect can interact with the pi-Diapason virtual machine in order to modify the service orchestration pi-Diapason code at runtime. Modifications are transmitted to the virtual machine and are then dynamically applied without interrupting the executing orchestration (the virtual machine supports services orchestration state consistency management). The service orchestration behaviour is changed on the fly by way of pi-calculus messages that contain the required changes.

SOAs defined using pi-Diapason can be checked against properties (eg deadlock-free, vivacity, liveness, structural and behavioural properties). Properties are defined using logic-Diapason, a logic-based properties definition lan-

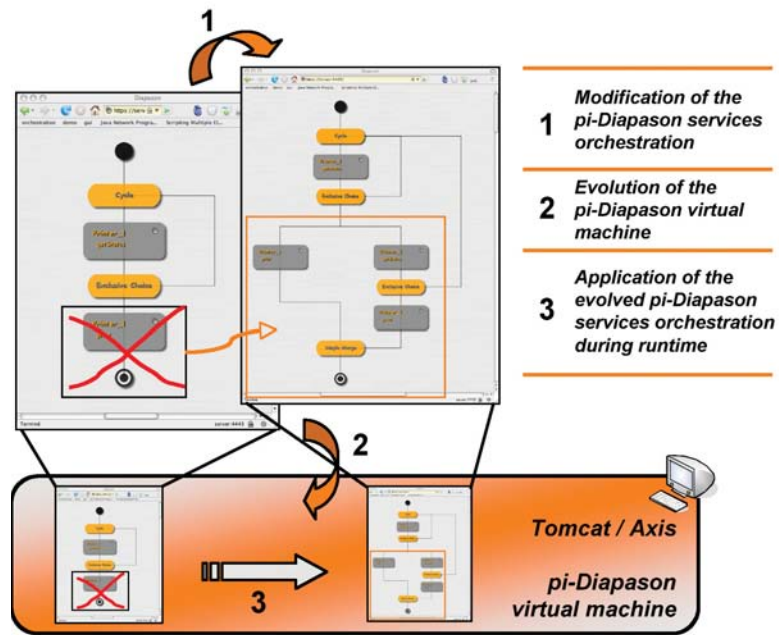


Figure 2: Orchestration evolution.

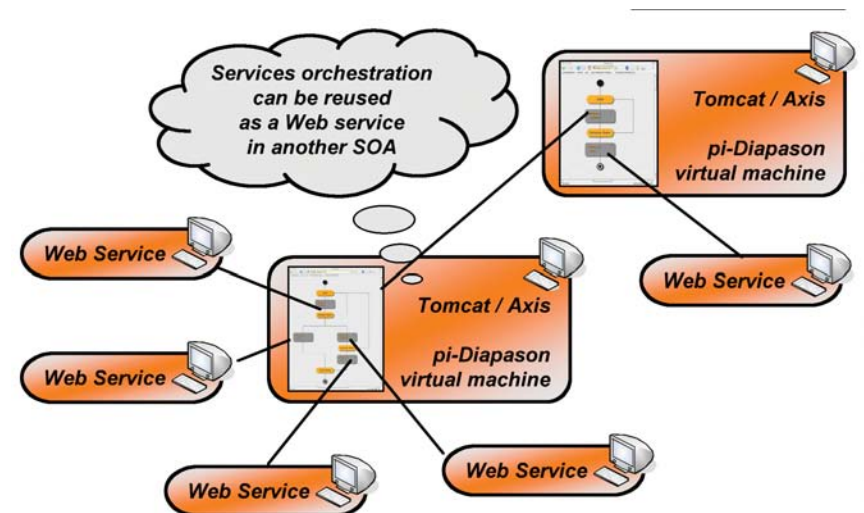


Figure 3: Orchestration reuse.

guage. SOA properties are then verified according to two steps: the first is performed by the pi-Diapason virtual machine that allows the extraction of all possible execution traces, while the second is done by the analyser, which allows properties expressed using logic-Diapason to be checked over all the previously extracted traces. In addition to this formal verification, the animator lets us simulate one or more of the

SOA's execution traces in a graphical manner. This is a more intuitive (though informal) way of checking the SOA's structure and behaviour.

Diapason Engineering Environment

A Diapason service orchestration is deployed as a service; this embeds the service-orchestration pi-Diapason description, a pi-Diapason virtual machine, in order to interpret the

orchestration and a context-aware deployment platform (according to the application server on which the service is deployed, ie Tomcat Axis or similar).

Please contact:

Frédéric Pourraz and Hervé Verjus
University of Savoie - Polytech'Savoie,
LISTIC, France
E-mail: frederic.pourraz@univ-savoie.fr,
herve.verjus@univ-savoie.fr

SUPER – Raising Business Process Management Back to the Business Level

by Matthias Born, Christian Drumm, Ivan Markovic and Ingo Weber

The SUPER project (Semantics Utilized for Process Management within and between Enterprises) is working on improving the modelling and management of business processes. This will be achieved by integrating semantic technology with business process management (BPM), and would provide answers to two of the most prominent issues emerging in this area: shifting control of processes from IT professionals to business experts, and scaling up BPM to support processes of higher complexity.

SUPER is an EU-funded Integrated Project that has a duration of three years and is under the coordination of SAP. It started in April 2006 and unites nineteen partners and approximately sixty researchers. The project consortium is a balanced blend of industrial partners, use case partners and academic research teams from all over Europe.

The motivation behind SUPER arises from the challenges of the increased frequency with which business models and the contexts of enterprises change in today's world. This rate of change is generally caused by new internal or external business requirements (eg closer integration with suppliers and customers, implementation of new industry-specific standards or deployment of new application components), which may

originate from emerging business opportunities or new regulations from legal bodies. In this environment, two requirements are paramount: to provide fast and easy access to the process space of an organization and to enable swift adaptation of operational business processes.

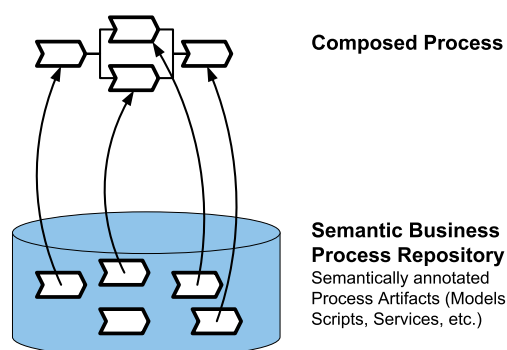
The major objective of SUPER is to raise business process management from the IT level to the business level. This objective requires that BPM be accessible to business experts and business analysts without requiring detailed technical expertise. Semantic Web and particularly Semantic Web services promise to enable users to perform complex tasks without requiring an understanding of the underlying technology. Therefore, this project aims at providing a framework that is context-aware and based on

Semantic Web Service technology, and which acquires, organizes, shares and uses the knowledge embedded in business processes and IT systems. This knowledge is presented to business experts and analysts in an understandable format through a novel process-modelling tool. Using this tool the framework enables them to easily analyse, change and create business processes, leading to a higher degree of agility in companies.

SUPER achieves this objective by adding semantic annotations to BPM artefacts (like process activities, services and execution artefacts), thereby making these artefacts accessible for advanced querying and reasoning. Using these querying and reasoning approaches, the tools developed in SUPER support users during business process modelling through techniques such as Semantic Business Process Discovery, Semantic Business Process Composition and Semantic Business Process Mediation.

Business Process Discovery provides support to the business expert during the modelling phase by simplifying the reuse of existing artefacts. For this task we provide a rich formal framework for the description of business process models, covering all workflow perspectives (data, control, resource). This enables

Figure 1: Business Process Composition based on semantic annotations of services and processes.



the business expert to pose expressive queries to the business process repository in the search for existing process components. Semantic Business Process Composition aims at enabling business experts to operationalize their business processes directly, by automatically deriving an executable process from a conceptual business process model (see Figure 1). Semantic Business Process Mediation facilitates Semantic Business Process Composition by enabling the seamless integration of processes originating from various stakeholders in a collaborative business process (see Figure 2). The use of such technology facilitates the task of modelling business processes in two ways: first, it improves the quality of the models through the reuse of established and optimized process components; and second, it

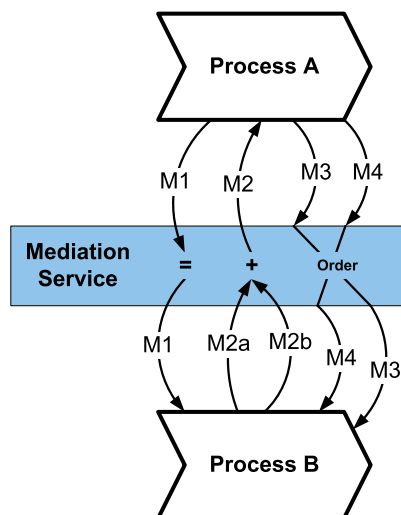


Figure 2: Business Process Mediation deals with heterogeneity in the behavioral interfaces and message formats of processes.

reduces the process modelling time by avoiding reinventing the wheel.

The chosen testbed for the achievements of the project is the telecommunication sector, where the limitations of traditional BPM approaches like cost and delay of process set-up are clearly visible, and currently limit the development of advanced, more agile business models.

Link:

<http://ip-super.org/>

Please contact:

Matthias Born, Christian Drumm,
Ivan Markovic and Ingo Weber
SAP Research, CEC Karlsruhe,
Germany

E-mail: {Mat.Born, Christian.Drumm,
Ivan.Markovic, Ingo.Weber}@sap.com

Life-Cycle Support of Semantic Web Services

by László Kovács, András Micsik and Tomás Pariente

In the frame of the INFRAWESBS project, a software toolset for creating, maintaining and executing WSMO (Web Service Modelling Ontology) services was developed. The implemented framework supports open and extensible development platforms for Semantic Web services.

The INFRAWESBS project has developed a framework that enables software and service providers to generate and establish open and extensible development platforms for Web service applications. The project was funded by the EU with participation of 12 European partners: FH Bochum, University of Innsbruck, IIT-BAS, SZTAKI, NTUA, Profium SA, Sirma SAI, FUTUREtec, Atos Origin, Best-HP, Aspasia Knowledge Systems and big7.net.

The project divides the life cycle of Semantic Web services (SWS) into two different phases: design time and run time. The role of the design-time tools is to ready existing Web services for semantic deployment. During the design-time phase, various tools and editors support the creation of semantic descriptions for existing Web services.

For example, the SWS Designer (SWS-D) is a visual editor of SWS descriptions, which helps to create service descriptions for Web services by applying drag-and-drop graphical editing and case-based reasoning. The Designer

contains a visual axiom editor, which makes the compilation of logical queries easier even for non-expert users.

The SWS Composer (SWS-C) is another visual tool for creating a Semantic Web service through the composition of existing WSMO-based services. It also uses a case-based memory for retrieving service composition templates quasi-similar to the service to be composed.

The resulting ontologies, goals and Semantic Web services are made accessible in a distributed registry (DSWS-R). WSML (Web Service Modelling Language) was chosen as the language for describing these semantic entities. The run-time phase involves discovery, selection and execution of the Semantic Web services. The Quality of Service (QoS) Broker component of the run-time environment collects QoS data for Semantic Web services, which are fed back into the phase of discovery and selection.

Developers of SWS-based applications use the SAM module (Service Access

Middleware) as a central API to INFRAWESBS services. With the help of SAM they are able to create Java code to discover, select and execute Semantic Web services. SAM is also able to hide the WSML syntax from developers, so no WSML parser is necessary in the applications.

The discovery engine of INFRAWESBS was implemented by SZTAKI as part of the SAM module. According to the usual scenario, the discovery engine receives a WSML goal as input and must provide a list of matching Semantic Web services, possibly coupled with additional information that supports ranking and selection.

Discovery implementation has three steps: pre-filtering, logical matching and finalizing to prepare the result. The aim of the pre-filtering step is to narrow the list of candidates using traditional text-processing (keyword matching) algorithms. This pre-filtering is supported by a structured text indexing and clustering service, which is another module of the framework. Pre-filtering is optimized

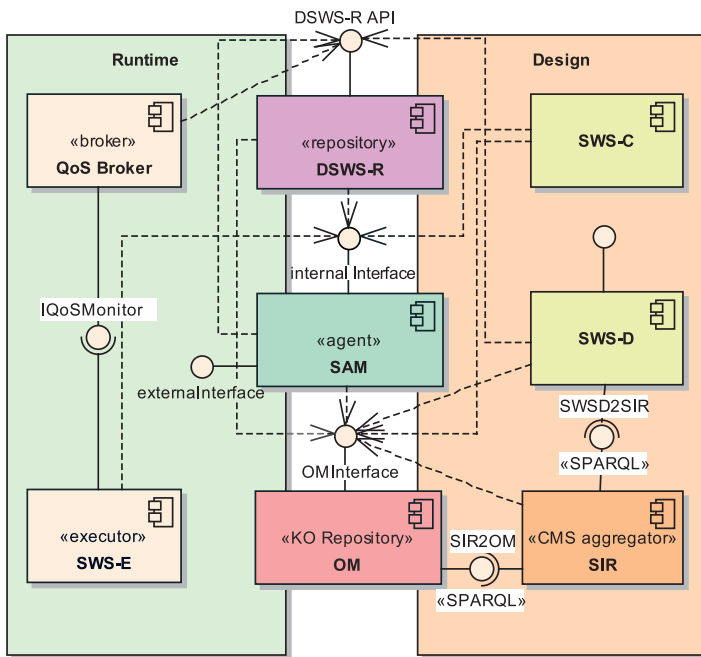


Figure 1: Overall architecture of INFRAWEBs modules.

for recall, so that it does not filter out possible matching candidates, yet it filters out definite non-matches using very fast textual queries.

Our implementation of semantic matching applies the unification facility of Prolog engines. Instead of matching large logical expressions, we break these expressions into small pieces using Disjunctive Normal Form (DNF), where the matching of each piece can be judged individually. This approach

also provides interesting possibilities such as the comparison, ranking and explanation of service matches (or non-matches). Furthermore, we are able to handle user preferences during discovery, which is becoming an increasingly important aspect of SWS usability.

In the final step of discovery, the list of matching services is enhanced with QoS data based on past execution experience, which can be used for service selection. QoS data is collected by the QoS Moni-

tor, a part of the Service Executor module. This Service Executor processes WSMO-based Semantic Web service descriptions using choreography and orchestration engines for executing specific rules, and performs the actual execution of the selected services.

The various tools implemented in the project are integrated by the INFRAWEBs framework, which is an enterprise service bus (ESB) connecting a peer-to-peer (P2P) network of nodes. Nodes can host Java and .NET services. Every peer can be deployed containing all or part of the INFRAWEBs stack of components. The framework hides from users the complexity of dealing with the P2P architecture.

On the other hand, the INFRAWEBs framework allows service application providers to create semantic-enabled applications based on the Semantic Web service descriptions stored in the P2P repositories. All methods provided by the components developed within the project are available through Java and Web service interfaces provided by the INFRAWEBs framework. They ensure interoperability and make the framework simple for developers to use.

INFRAWEBs reflects a novel approach to solving problems that occur during the creation and maintenance of Semantic Web services and their applications. It is based on tight integration of similarity-based and logic-based reasoning. Similarity-based reasoning is used for fast finding of approximate solutions, which are further concretized by the logic-based reasoning.

Most of the tools and the integration framework are available as open-source software at the project Web page, under LGPL license.

Link:

<http://www.infrawebs.eu/>

Please contact:

Tomás Pariente
 ATOS Research & Innovation, Spain
 Tel: +34 91 214 9321
 E-mail:
 tomas.parietelobo@atosorigin.com

László Kovács

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

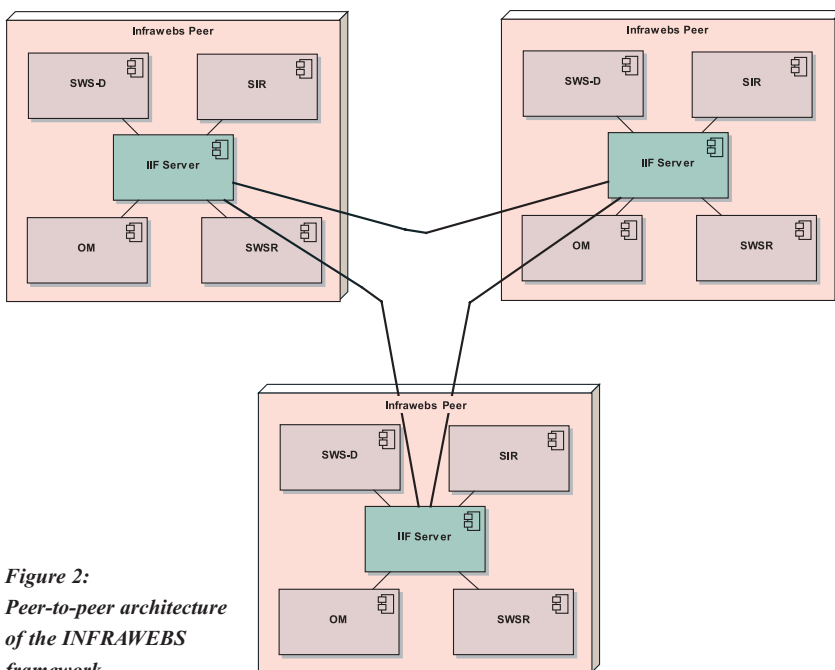


Figure 2: Peer-to-peer architecture of the INFRAWEBs framework.

Aligning Virtual Organizations with SOA

by Mikko Salonen and Jyrki Haajanen

Virtual Organizations (VOs) are collaborative groups that can be formed in several ways, and for several reasons. Typically they are logical entities, have a limited lifetime, are geographically dispersed, and are created to solve a specific problem or to enhance and develop business processes. Information networks play a significant role in the interaction of VOs, and their importance is expected to grow even further. Service-oriented architecture (SOA) is an architectural approach based on the service-oriented computing (SOC) paradigm. Since SOA encompasses architecture issues in both business and information systems, it has great potential to enable alignment of VO approaches with IT-based business networking approaches, such as enterprise interoperability. So while SOA can contribute to VOs, VOs can also make their own demands on and contributions to the adaptation of the new architecture.

This research is part of a project named SOAMeS (Service-Oriented Architecture in Multichannel e-Services), at VTT Technical Research Centre of Finland. The project, which commenced in mid-2006 and will last until the end of 2007, is running in collaboration with the University of Helsinki. The main funding bodies are Tekes (the Finnish Funding Agency for Technology and Innovation) and VTT, while four Finnish companies also provide funds and are involved in its execution. The project aims to clarify the potential of service-oriented solutions in business network management and strategic planning. Furthermore, the project explores the usability of existing SOA tools for multi-channel service composition, and elaborates those tools by enhancing collaboration modelling with dynamic and non-functional aspects. The aim of the VO research is to study what SOA can offer to VOs, and to identify the challenges it may confront, the demands VOs will set for SOA, and the contribution that the VO approach can provide to other SOA-enabled business networking approaches.

Business environments are becoming more and more dynamic, and information systems need to keep up with this pace. New architectures must be able to adjust to and even enhance the ongoing change and development of business environments and organizations. It is important that the service-oriented architect involved in development understands the issues, demands and requirements – both technical and business-related – from both sides.

Virtual organizations need tools for, among other things, preparatory specification, consortium formation, support

and finalization. Different actors, for instance, may have different operating systems, data formats and languages. A standard or architecture must therefore be adopted that can integrate systems and embrace requisite services. Innovation, different communication types, data sharing, decision making and other synchronous and asynchronous collaboration and sharing methods are necessary for virtual organizations. Tools must support the processes and in doing so help improve interaction within VOs. Dynamic operation requires that ICT solutions are flexible and agile, and can enable changes in relationships within and between virtual organizations.

Given the organization-specific requirements listed above, SOA seems to be the natural choice for solving the problems resulting from organization-specific information systems and business processes. SOA provides a natural isolating layer between the business requirements and their implementation, while still allowing their direct linkage in implementation. Full utilization of SOA principles with nested services results in a more economical reuse-oriented environment, where a single business process can be expressed as a service consisting of subservices. These subservices (eg customer information updates) can be reused, thereby improving data and system integrity, focus of development costs and architectural control.

The main challenge in applying SOA is, however, of a psychological kind. SOA is more a management philosophy or discipline than a plain information system architecture, and thus should be applied in a determined way. Managers

should not expect massive returns immediately, but should be prepared to wait for a sufficient period of time for them. This does not mean, of course, that applying SOA will not bring immediate results, but rather that over a short time they are typically exaggerated.

The aim of this research is to find the link between SOA and VOs, to discover how SOA can be used in VOs, and also to identify whether VOs can contribute something new to SOA. How can they be combined so that the best possible result will be achieved, especially from the point of view of the organization's management and strategic development? The role of virtual organization is growing and SOA may enable its better use and collaboration.

Link:

<http://soames.cs.helsinki.fi/>

Please contact:

Mikko Salonen
VTT, Technical Research Centre
of Finland
Tel: +358207224194
E-mail: mikko.salonen@vtt.fi

Jyrki Haajanen
VTT, Technical Research Centre
of Finland
Tel: +358207226028
E-mail: jyrki.haajanen@vtt.fi

A Pervasive, Service-Oriented Architecture for Supporting Teamwork

by Schahram Dustdar and Hong-Linh Truong

At the heart of the EU 'inContext' project, the Pervasive Collaboration Services Architecture (PCSA) aims at providing a pervasive, SOA-based architecture for supporting various kinds of teamwork. This architecture comprises different kinds of Web services, loosely coupled in a dynamic environment that includes diverse underlying operating systems and networks, necessary for collaboration and teamwork. The goal of this platform is to reduce as far as possible human intervention in the support of collaborative work, by means of autonomic capabilities based on context information and interaction patterns.

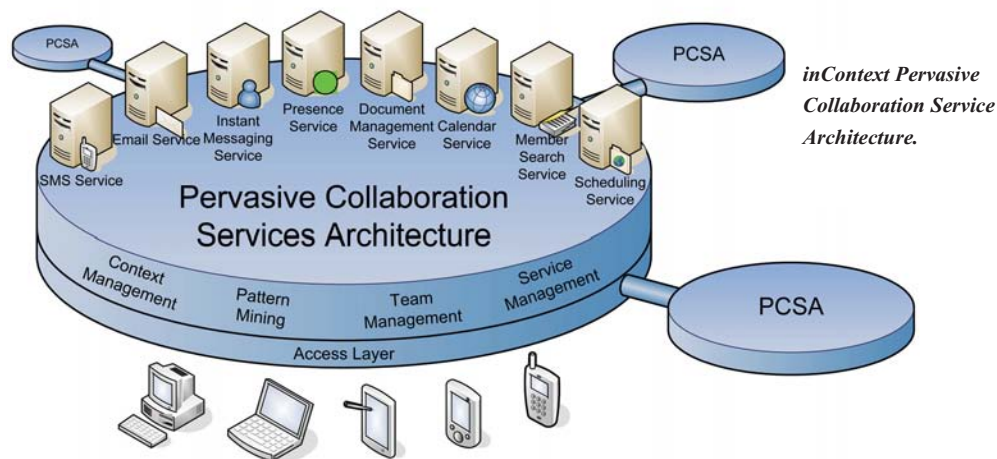
Recent advances in computing devices and network technology are creating diverse, pervasive environments and opening various opportunities for human collaborative work on these environments. Nowadays, members of teams in collaborative processes may span different times and spaces. Depending on the context of a collaboration, a team may be set up through an Internet-based or mobile ad hoc infrastructure, using mobile devices and/or high-end computers, working together for only a short period of time or for up to several days or weeks, just to name a few cases. More than ever before, team-

The inContext Project

The inContext project aims at supporting highly dynamic forms of human collaboration such as Nimble (short-lived collaboration to solve emerging problems), Virtual (spanning different geographical places and involving diverse professionals) and Mobile (collaboration with mobility capabilities) teams.

These teams require different mechanisms for coordination, and in many cases also different services (eg document sharing, project management and instant messaging) and infrastructures

As shown in Figure 1, this architecture comprises collaboration services and the inContext platform. In order to support basic team activities, we need plenty of common collaboration services, such as calendars, task management, instant messaging and document management. However, a variety of other services need to be developed for supporting emerging teamwork and autonomic capabilities. Hence, on the one hand we adopt and integrate existing common services by wrapping them, and by building and extending their Web service interfaces since not all services provide one. On the other



work is highly dynamic and flexible in terms of both time and space; existing infrastructures in which teamwork is supported by fixed, tightly coupled systems that include dedicated services and portals but which do not interact with each other, thus fail to provide sufficient capabilities. Not only will an SOA-based platform allow us to integrate various existing collaboration services into dynamic, pervasive environments, but newly developed collaboration services for emerging teamwork could also be easily plugged into such environments.

(eg large-scale and Internet-based mobile devices, and mobile ad-hoc/P2P networks). SOA-based solutions thus offer greater advantages for inContext over other solutions, such as those that are portal-based.

The inContext PCSA

The inContext PCSA (Pervasive Collaboration Services Architecture) aggregates different types of services, all based on Web Services, to support collaborative team processes. These services are loosely coupled and can be deployed in various hosting environments.

hand, services in the inContext platform are able to manage context information, analyse interaction patterns and perform autonomic capabilities for the PCSA. This means the PCSA is able to fulfil the needs of different teams and to cope with changes in team forms as they evolve.

The PCSA Network

Since team collaboration is conducted in a distributed system involving multiple organizations and spanning multiple locations, we have to ensure that team members are able to access all the serv-

ices in a transparent way. An efficient way to support these collaboration scenarios is to connect multiple PCSAs in a peer-to-peer network through which a team member can perform his/her collaboration without needing to know where and how services are deployed. As depicted in Figure 1, the inContext platform and in particular the Access Layer will route invocations to appropriate services in the PCSA network. In this respect, context information plays a key factor as it provides important information about the location and status of collaboration services.

Autonomic Collaboration Architecture

The inContext PCSA integrates and provides various services for human collaboration. Furthermore, by offering autonomic capabilities, it aims to reduce human intervention in managing the collaboration. Being autonomic requires an intensive use of context information and interaction patterns to

adapt and provision services to the changes in and requirements of teams and their respective environments. Take the scheduling of a meeting as an example in which autonomic capabilities are important. In normal environments, the initiator of the meeting must select members based on roles and meeting objectives, determine possible dates based on their calendar, select potential days, send the selected days to members via e-mails or instant message or SMS, search for relevant documents needed for the meeting, create an agenda and so on. All of these steps are performed manually. With support of the inContext PCSA, this level of human intervention can be reduced. For example, the initiator may need to simply select a checkpoint in the project timeline and ask the PCSA to plan a meeting.

Based on context information about members' presence and location, human intervention in many steps like

selecting members, checking calendars and sending confirmations can be completely removed. In the inContext project, autonomic aspects will be focused on work activities and the required service adaptation.

The inContext research is coordinated by the Vienna University of Technology and is conducted together with Softeco Sismat SpA, DERI Ireland, European Microsoft Innovation Center, Electrolux Home Products Italy SpA, Hewlett Packard Italiana, the University of Leicester, West Midlands LGA and COMVERSE Ltd.

Link:

<http://www.in-context.eu/>

Please contact:

Schahram Dustdar

Vienna University of Technology,
Austria

Tel: +43 1 58801 18414

E-mail: dustdar@infosys.tuwien.ac.at

Engineering Services

by Mike P. Papazoglou

Service-oriented computing is not simply about deploying software: it also requires that organizations evaluate their business models and come up with service-oriented design and development techniques and support plans. At INFOLAB/University of Tilburg in the Netherlands we have developed an experimental methodology for service-oriented design and development that applies equally well to Web services and business processes.

Service orientation utilizes services as constructs to support the rapid, low-cost and easy composition of distributed applications. Key to this concept is the service-oriented architecture (SOA), which is a logical way of designing a software system to provide services to either end-user applications or to other services distributed over a network, via published and discoverable interfaces. A well-constructed SOA can empower a business environment with a flexible infrastructure and processing environment by provisioning independent, reusable automated business processes (as services) and providing a robust foundation for leveraging these services.

In their early use of SOA, many researchers and developers think that they can port existing components to act as Web services just by creating wrappers and leaving the underlying compo-

nent untouched. Since component methodologies focus on the interface, many developers assume that these methodologies apply equally well to SOAs. Thus, the placement of a thin SOAP/WSDL/UDDI veneer on top of existing applications or components that implement the Web services is now widely practised in the software industry. Yet this is in no way sufficient to construct commercial-strength enterprise applications. Unless the nature of the component means it is suitable for use as a Web service – and most are not – it takes serious thought and redesign to properly deliver a component's functionality through a Web service.

While relatively simple Web services may be built with conventional development methodologies, a service-oriented development methodology is of critical importance to specify, construct, refine

and customize highly volatile business processes from internally and externally available Web services.

Our service-oriented design and development (SoDD) methodology provides sufficient principles and guidelines to specify and construct business processes choreographed from a set of internal and external Web services. It takes into account a set of development models (eg top-down, bottom-up and meet-in-the-middle), stresses reliance on reference models, and considers several service realization scenarios, including green-field development, outsourcing, and legacy wrapping in cases where services are assembled out of pre-existing components.

Our service-oriented design and development methodology is based on an iterative and incremental process that com-

prises one preparatory and eight distinct main phases that concentrate on business processes and may be traversed iteratively. These are planning, service-oriented analysis and design, construction and testing, provisioning, deployment, execution and monitoring.

Service-Oriented Analysis and Design Phases

Service analysis aims at identifying, conceptualizing and rationalizing business processes as a set of interacting Web services. In particular, the analysis phase places an emphasis on identifying and describing the processes and services in a business problem domain, and on discovering potential overlaps and discrepancies between processes under construction and available system resources that are needed to realize singular Web services and business processes. It therefore examines the existing services portfolio at the service provider's side to understand which patterns are in place and which need to be introduced and implemented. Service analysis helps prioritize business processes and services where SOA can contribute to improvements and offer business value potential. It also helps to centre efforts on business domains within an enterprise that can be mapped to core business processes.

The analysis phase reviews the business goals and objectives of an enterprise, since these drive the development of business processes. It helps focus SOA initiatives by creating a high-level process map that identifies business domains and processes of particular interest to an enterprise. Business processes are ranked by criteria related to their value and impact, reuse and high consumption, feasibility and technical viability. From the process map, analysts can identify candidate business services that relate to these processes. Candidate business services are those that have potential value for an organization and can be evaluated on the basis of reuse, business impact, and organizational value.

Designing a service-oriented application requires developers to define related, well-documented interfaces for all conceptual services identified by the analysis phase, prior to constructing them.

The design phase encompasses the steps of singular service specification, business process specification, and policy specification for both singular services and business processes. Service design is based on a twin-track

design approach that provides two production lines – one along the logical part and one along the physical part of the SOA – and considers both functional and non-functional service characteristics. The purpose of logical service design is to define singular services and assemble (compose) services out of reusable singular service constellations. This calls for a business process model that forces developers to determine how services combine and interact jointly to produce higher level services. The physical design trajectory focuses on how to design component implementations that provide services at an acceptable level of granularity. Physical design is thus based on techniques for leveraging legacy applications and component-based development.

Link:

<http://www.tilburguniversity.nl/faculties/feb/organisation/dept/ism/research/themes/soc/>

Please contact:

Mike P. Papazoglou
University of Tilburg, The Netherlands
E-mail: mikep@uvt.nl

Business Process-Driven Service Architecture

by Veronica Gacitua and Claus Pahl

Increasingly, enterprises are using service-oriented architecture (SOA) for software application integration. The technical problem – integration at the service level – is one of architecture. The business model domain that drives the integration can, however, be linked to the service architecture level by adding a service-centric architectural perspective into business process modelling. Service architectures can be reused in the form of architectural patterns and styles, which provides quality improvement and cost reduction.

IT Architecture and Enterprise Application Integration

An IT architecture is an enterprise-wide system of software applications. Enterprise application integration (EAI) methodologies provide integration solutions based on domain and application engineering techniques. In conjunction with SOA, the architecture focus is evident and forms the core of this project.

Reuse of Business Process-Driven Service Architecture

Focusing on architectures can bridge the gap between business domain-level con-

cerns and service-level platform aspects. We propose a three-layered framework with business domain modelling, application architecture and service-centric integration architecture to structure an integration problem:

- at the business modelling layer, the business processes constrain and drive the integration of software applications
- at the application architecture layer, business process models are refined into architecture design models by means of architectural constraints and service classification and identification

- at the integration architecture layer, the actual integration of software applications is realized in the form of service compositions.

Our project objective is to enhance this incremental process by adding a reuse focus that allows architectures to be reused in the form of enterprise-wide architectural patterns and styles. This pattern-based approach serves to guide SOA-based application integration and migration. A coherent, service-centric architecture-modelling notation based on the recently stan-

standardized Business Process Modelling Notation (BPMN) is the central element of this integration architecture approach. The reuse aim is addressed through style-based and pattern-based modelling and a pattern refinement calculus.

This project is carried out by the Software and Systems Engineering Group at Dublin City University. The project was motivated as a result of our collaboration with major software consultancy firms and the involvement of team members in large-scale SOA and EAI projects in sectors such as mining and IT providers, where questions of quality and cost-effectiveness and the possibility of reuse to address these questions have emerged.

Process and Service Architecture

The first objective of our framework is software architecture-driven integration. An incremental architecture development technique is our proposed solution. Software architecture is the instrument with which to realize the integration of heterogeneous applications in an IT architecture.

BPMN provides the notational basis for our architecture-driven integration. We extend BPMN to a service-centric architecture description language to enable the enrichment of business models towards a service-centric business architecture. Business domain models provide a complementary information architecture on top of business processes. Together, they can be constrained and refined through reference architectures and architectural styles at the service-centric application architecture level. A graph-based formalization of BPMN allows rigorously defined architecture constraints to be used.

Pattern-Driven Architecture Reuse

The pattern-driven reuse of service architectures is the second framework objective. Patterns are used in software design to structure architectures and allow the reuse of successfully applied architectural designs. Patterns are architectural abstractions that can cross and therefore integrate existing service application boundaries.

We focus on architectural patterns in terms of process activities and services

and their part in the composition of processes. The BPMN architectural extension is the basis for pattern-based techniques:

- pattern identification is supported by a BPMN-based pattern language for both business-level and service-level patterns, at the business domain layer and at the integration architecture layer respectively, and an empirically developed repository of patterns
- pattern mapping – transforming incrementally from process patterns to architectural service-based patterns using application integration layer techniques – addresses architectural style and pattern refinement, supported through a pattern refinement calculus.

Graph-theoretic formalisms define the service-based pattern language and calculus.

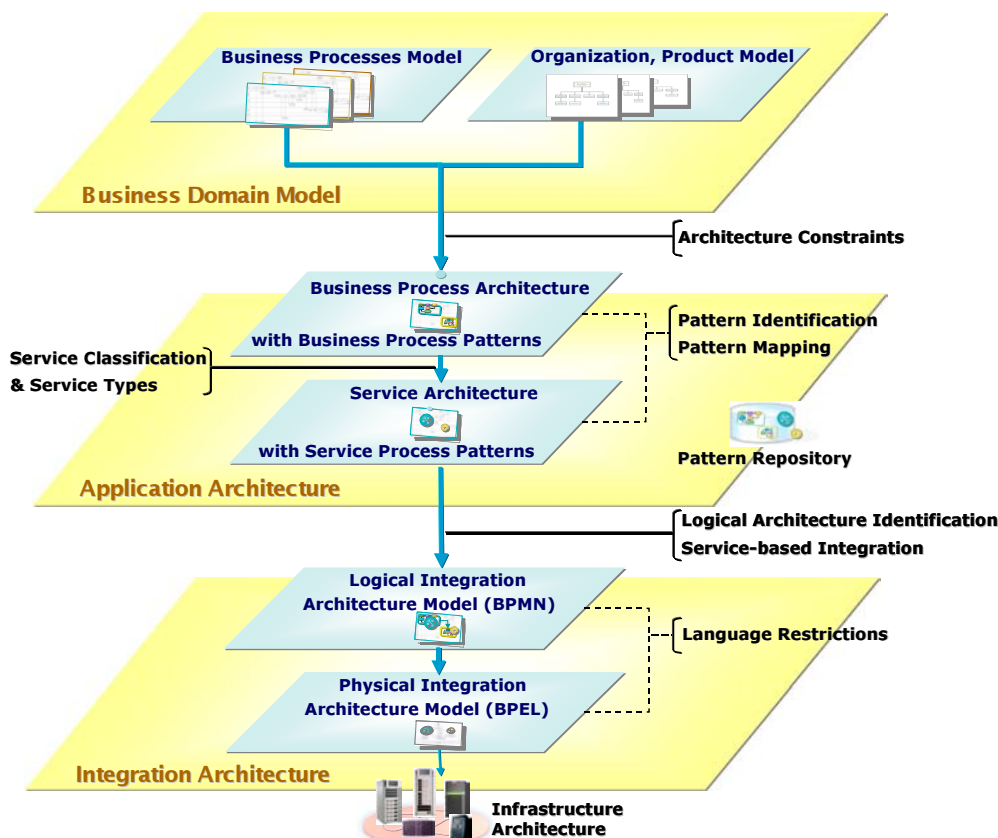
Conclusions

Bringing quality and cost-effectiveness into service-based IT architecture integration and migration through reuse are the benefits of the proposed integration architecture approach. Techniques such as architectural patterns have quality attributes associated to them, like performance or reliability, which capture observable properties of software applications. The software architect can influence these application system qualities based on a suitable choice of architectural patterns already at the service level. Other, more design-level quality attributes, such as maintainability, are also strongly affected by the application's architecture. Quality-driven architecture is the essence of the proposed approach.

Model-driven development for architecture-based integration is another perspective that explains the benefits of this approach in terms of layered abstraction up to the business level and at least semi-automated incremental transformation and refinement of service architecture descriptions.

Please contact:

Claus Pahl
School of Computing, Dublin City
University
Tel: +353 1 700 5620
E-mail: Claus.Pahl@dcu.ie



Layered Business Process-Driven Service Architecture.

A Service-Oriented Middleware for B3G Networking

by Valerie Issarny, Mauro Caporuscio, Pierre-Guillaume Raverdy

Beyond 3rd Generation (B3G) networking provides mobile users with unique features for seamlessly accessing networked services. However, the provision of services over B3G distributed computing platforms faces numerous challenges, and these are being investigated by the INRIA ARLES project team.

The ARLES (Software Architectures and Distributed Systems) project team at INRIA is investigating solutions to issues that occur as part of the design and development of a service-oriented middleware for B3G networking. These include: (i) developing services that can be easily deployed on a wide range of evolving infrastructures, from networks of devices to stand-alone wireless resource-constrained handheld devices; (ii) making services context-aware so that they can benefit from networked resources and related services in the most effective way; and (iii) ensuring that users always experience the best possible quality of service according to their specific situation.

B3G networking offers broad connectivity through the various network technologies that are now becoming available to the vast majority of end-users, thanks to the newest multi-radio devices like smart phones embedding, eg UMTS, WiFi and Bluetooth networking (see Figure 1). Still, the network's diversity and richness must be comprehensively exploited by the software applications made available to nomadic users. Service-oriented architecture (SOA) appears to be an appropriate paradigm for engineering software applications to be deployed in the B3G network, as functionalities provided by networked resources may be conveniently abstracted as services. Specifically, a service represents an autonomous networked entity that provides a set of functionalities to its environment, while the latter continuously undergoes alterations, due in particular to changes in network connectivity. Thanks to service discovery protocols, networked services may be dynamically located for use by service clients, as hosts of clients and services move into communication range of each other (either via the network infrastructure or directly in an ad hoc way). Furthermore, networked services may be

deployed on the various types of nodes composing the B3G network, from nodes of the core network infrastructure to end-user nodes. Such a feature allows rich networked software services to be developed and deployed by various actors, since there is no longer dependence on the core network infrastructure.

As part of the European project 'PLASTIC', the INRIA ARLES project team is designing and developing middleware to support the deployment and dynamic composition of mobile, adapt-

that are now embedded on wireless devices should be comprehensively managed by the middleware, offering the abstraction of an integrated multi-radio interface to the software services of the upper layers.

The aforementioned focus on the middleware leads directly to the high-level architecture depicted in Figure 2. The middleware architecture is layered on top of a legacy networked software platform, and decomposes into two main layers to enable distributed applications in the B3G network. Applica-



Figure 1: B3G environment.

able services in B3G networks. More specifically, we are investigating solutions to B3G networking at the middleware layer, with a special focus on enabling the deployment of services over diverse terminals, including those that are (mobile) wireless and resource-constrained. The middleware will enable the effective exploitation of B3G networking capabilities, which include composing the various networks in reach to improve availability of services and further offering seamless mobility. In addition, the various radio interfaces

include Web services or any other application software that uses (a subset of) the middleware API.

The lower communication middleware layer deals with service provisioning in the B3G network, which includes enabling the execution of services on the end-users' wireless handheld devices, which embed various radio network interfaces. This layer offers the abstraction of an integrated multi-radio network, which comprehensively composes any networks within reach via the

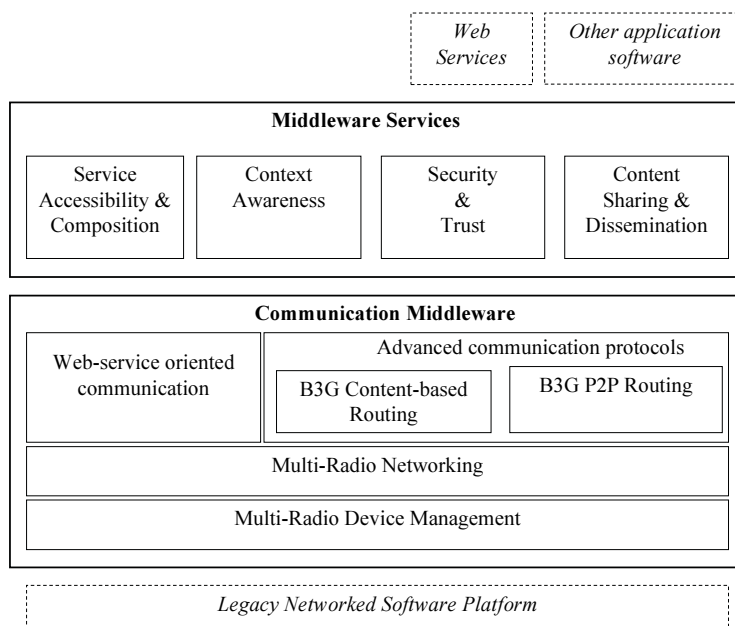


Figure 2:
Middleware architecture.

embedded radio interfaces (eg GPRS, WiFi, Bluetooth), as identified by the multi-radio device management functionality. Basically, the multi-radio networking layer offers point-to-point and multicast messaging over the multi-radio network, and chooses the underlying network over which messages are to be actually sent. This is done according to application-layer requirements in terms of both functional and non-functional properties. Advanced routing protocols are also offered on top of the multi-radio networking layer.

Communication based on Web services communication enriches the traditional functionalities of a SOAP engine to allow for SOAP-based interaction in the B3G network. This includes (i) enabling access to services that may be in distinct networks thanks to multi-network routing, and (ii) dealing with seamless mobility as long as the respective hosts of the given service's client and provider remain within reach via at least one radio link of the B3G multi-radio networking environment. The PLASTIC project also investigates additional advanced communication protocols that are customized to the open B3G networking environment, such as content-based and P2P routing.

The upper middleware services layer embeds advanced services related to distributed resource management. Service accessibility and composition enable (i)

dynamic service discovery in the changing networking environment, and (ii) distributed services through the composition of networked services. Context awareness is crucial to service adaptiveness, since it involves aggregating relevant context information sensed in the environment and making it available to application services for application-specific adaptation. Security and trust are prime requirements for service provisioning in an open, wireless computing environment. Content sharing and dissemination are also major middleware functionalities, and need to be tuned to the specifics of B3G networking.

The proposed service-oriented middleware has been designed to enable the deployment of services on B3G-enabled mobile, resource-constrained devices, thus exploiting their multi-radio networking capabilities. We are now refining the architecture as part of the prototype implementation of the various modules of the middleware. Still, some functionalities require further research due to the inherent complexity of the B3G networking environment. Thus, the prototype implementation of the middleware will be enhanced gradually, with advanced research solutions being substituted for more basic solutions that were elicited to enable early experiment/validation.

Our research is conducted as part of the European IST PLASTIC project.

Links:

ARLES: <http://www-rocq.inria.fr/arles/>
PLASTIC: <http://www.ist-plastic.org>

Please contact:

Valerie Issarny
INRIA-Rocquencourt
Tel: +33 1 39 63 57 17
E-mail: Valerie.Issarny@inria.fr

Exploiting Organizational Models for Semantic Service Description, Matchmaking and Composition in Service-Oriented Multi-Agent Systems

by Alberto Fernández and Sascha Ossowski

In open multi-agent systems (MAS), agents are often conceived as software entities capable of advertizing the services they provide, locating other service providers that offer services potentially of interest to them, and negotiating agreements regarding service enactment. Although organizational models are usually present in agent-oriented design methodologies, they have not yet found their way into service description mechanisms for agents and consequently are not explicitly exploited by today's service discovery and composition mechanisms. We propose an approach that bridges this gap by making use of organizational concepts such as roles and interactions to extend current trends in semantic service description, matchmaking and composition in MAS.

In service-oriented MAS, agents need to be more than mere wrappers of the service interface: besides executing a service, they should be able to engage in different types of social interactions when providing services. In a health-care assistance scenario, for instance, it is not enough for an agent providing a second-opinion service to come up with a diagnosis based on the service's input data; it should also be able to explain the outcome of service execution, recommend a treatment, or even query the client for more details about symptoms.

We have developed an organizational modelling technique that captures the different types of social interactions that can occur during the provision of a service, as well as the roles that take part in those interactions. The result is a taxonomy of roles and interactions that we use for service description, matchmaking and composition.

Service Description

Agents use a service description language to specify relevant characteristics of the service(s) they provide. Most current approaches base their service descriptions on inputs, outputs, preconditions and effects. Sometimes additional parameters, like 'category', are taken into account. In addition to this, our approach exploits organizational concepts to further characterize the contexts in which services can be used. We distinguish between service advertisements and service requests.

In service advertisements we specify (i) the role played by the provider in the interaction, and (ii) a set of roles that must be played by the requester agent for the correct accomplishment of the

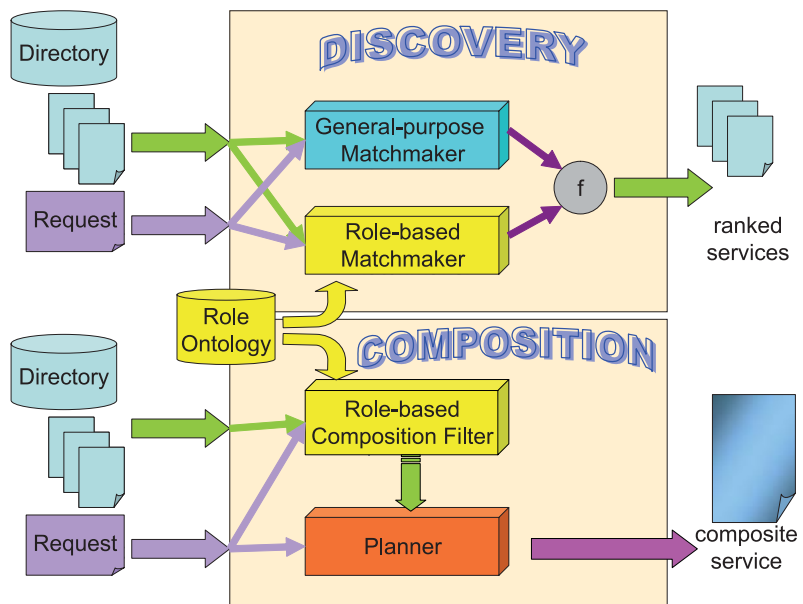


Figure 1: Service discovery and composition processes.

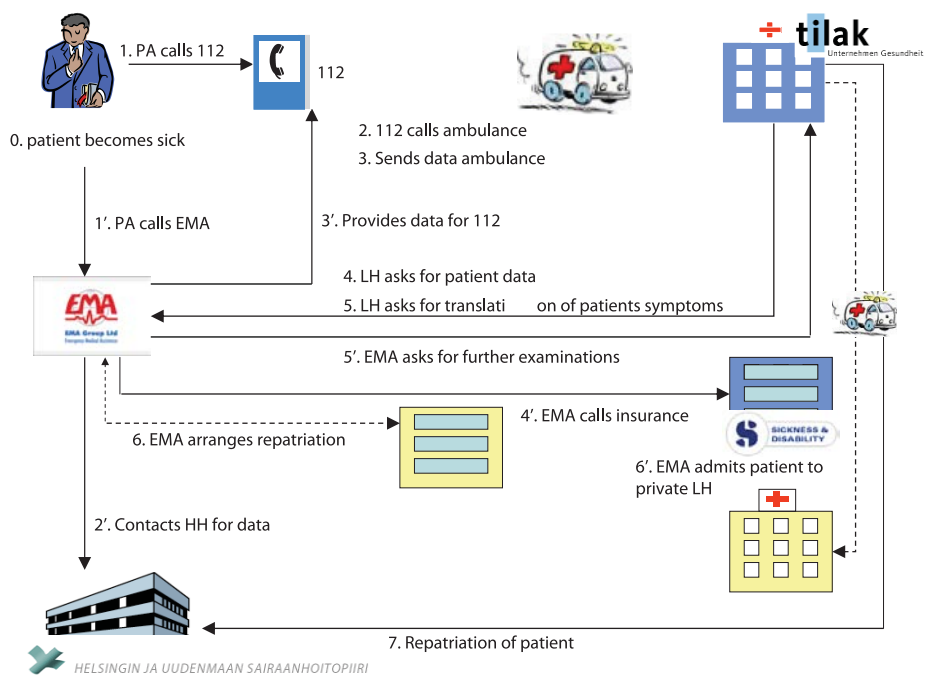


Figure 2: CASCOM emergency assistance application scenario.

service. For example, a second-opinion provider may require that the requester be capable of providing information about their symptoms.

Service requests comprise two elements: (i) the desired roles of the service provider, and (ii) the set of roles that define the capabilities of the requester. For example, the requester of a second opinion can specify that it is able to provide information if needed.

Service Matching

We have developed a role-based matching algorithm that takes as input a service request and a service advertisement, and returns the degree of match between them. This is a real number between 0 and 1 by which service providers can be ranked for selection.

The semantic match of two roles, R_A (advertisement) and R_Q (query), is made based on the ontology of roles. It is a function that depends on the level of match (the subsumption relation between the two roles in the ontology), and the distance (number of arcs) between R_A and R_Q in the taxonomy. The semantic match between two services is made by searching for the role in the advertisement that best matches that in the query. Minimum and maximum

are used for aggregation in the case of conjunctions and disjunctions in expressions, respectively.

Service Composition

If no adequate services are available for a specific request, a planning functionality can be used to build up composite services. In open large-scale service-oriented MAS, a pure AI planning approach can become impracticable due to the vast number of services (operators) that are usually registered in the directory.

We have put forward a method that exploits the MAS organizational information in order to heuristically filter out those services that are probably irrelevant to the planning process. Our heuristic is based on the plan dimension and on the number of occurrences of services in plans: the more important a service is, the greater the number of plans for which it is necessary and the shorter the plans for which it is required. We can approximate this information by storing and processing the plans historically created. However, the number of services and possible queries may become too large and the continual repetition of one particular service request is rather unlikely. To overcome this drawback, we cluster

services into classes based on certain properties. In particular, we obtain service class information from the role and interaction taxonomies that are derived from the MAS organizational model.

Application

The approach presented here has been implemented as part of the European IST project CASCOCOM. This project is developing, implementing and validating an agent-based service coordination infrastructure for the discovery, composition and execution of innovative Semantic Web services across mobile and fixed P2P service networks. In particular, as a component of the infrastructure, it is part of a field trial in medical emergency assistance that is currently running at the Tyrolean Hospital Consortium in Austria.

Link:

CASCOCOM: <http://www.ist-cascom.org/>

Please contact:

Alberto Fernández

University Rey Juan Carlos,
Madrid, Spain

Tel: +34 91 488 7084

E-mail: alberto.fernandez@urjc.es

Building Service-Oriented Web Applications for Business Processes

by Victoria Torres, Vicente Pelechano and Pau Giner

Business Processes (BPs) play a very important role in the development process of applications. They allow us to specify organization goals by means of tasks and the participants in charge of these tasks. Moreover, taking into account that the Web is being established as a software development platform, it is necessary to provide methods and tools that allow the systematic construction of these applications. This work presents a method for constructing service-oriented Web applications that provide not just data management, but also support for distributed business process execution.

Web applications can no longer be conceived as isolated systems. In fact, in order to achieve sets of shared goals in a business environment, applications must cooperate with each other. The increasing adoption of Web services technology facilitates this cooperation, and this work presents an extension to a Web Engineering method for the systematic construction of business-driven Web applications based on a service-oriented paradigm.

This extension involves reconsidering the method in use at two different levels. On the one hand we have extended the method at the modelling level, firstly introducing the required primitives in both the navigational and presentation models to cope with BP issues, and secondly introducing the Business Process Model (BPM) for the specification of BPs (see Figure 1). On the other hand, at the architectural level we have introduced a process engine (see Figure 2),

which is in charge of orchestrating the set of services that form the modelled BPs.

As Figure 1 shows, the method comprises a set of models that allow us to specify, in the problem space (the modelling level), the system being developed. Briefly, these models capture system structure, behaviour, navigation and presentation. In addition, the BPM was introduced into the method. This allows

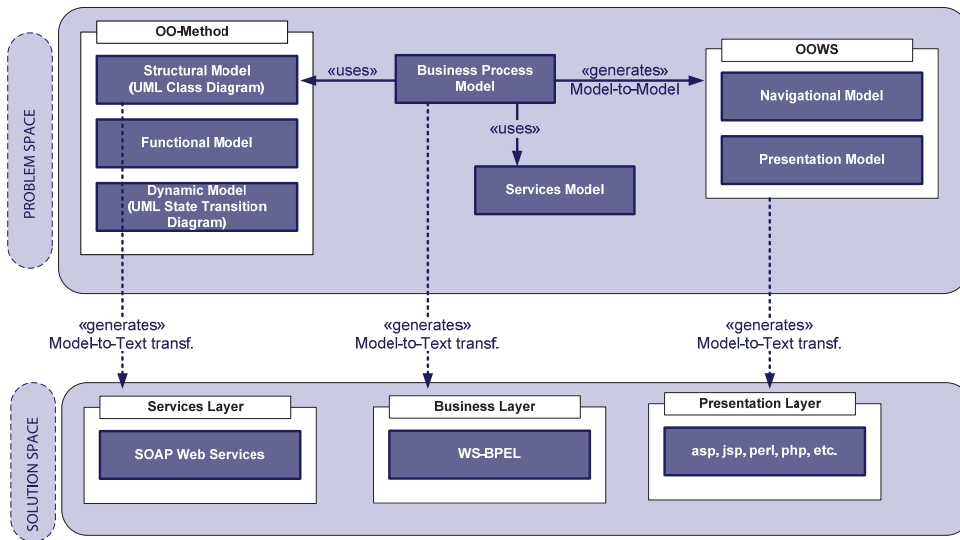


Figure 1: Extension defined at the modelling level.

us to specify which BPs must be supported by the Web application. These BPs are specified as a set of organized tasks that are performed either by the functionality defined in the Structural Model or by functionality that is included in the Services Model and is provided by external partners. The notation used to specify this model is the Business Process Modelling Notation (BPMN) OMG Specification.

Moreover, to build the model-equivalent artefacts for implementation we rely on model transformation techniques. We have defined both Model-to-Model and Model-to-Text transformations, which we use depending on the intended target artefact. Model-to-Model transformations generate the Navigational Model (NM) that supports the corresponding BPs. Model-to-Text transformations translate (i) NMs into Web pages implemented in a specific

language (asp, php, perl, jsp etc), and (ii) BPMs into a BP-executable language (WS-BPEL, WSFL, XLANG, BPML, Petri Nets etc).

Regarding architectural extensions, we have introduced two new elements at the business logic level (see Figure 2). The first element refers to a process engine. This element is in charge of orchestrating different services that join together to accomplish a specific goal. In particular, due to increasing adoption of Web service technology, we have considered the introduction of process engines that provide support for the Web Services Business Process Execution Language (WS-BPEL). This language allows us to describe the behaviour of a BP based on interactions between the process and its partners through Web service interfaces. However, as this language was not designed to support processes that involve inter-

action with people, we had to introduce a second element into the architecture. This element refers to the Task Manager (see Figure 2) and behaves as a mediator between the process engine and the Web application when human tasks are invoked.

As a result, we are combining Web Engineering best practices with Business Process Management solutions in order to endow Web applications with support for BP execution. The Web Engineering community can rely on existing process engines (which control message interchange in an organized fashion) to build Web applications that provide support for process execution. At the same time, BPM solutions can be endowed with richer graphical user interfaces that are defined at the modelling level and are completely integrated with Web applications.

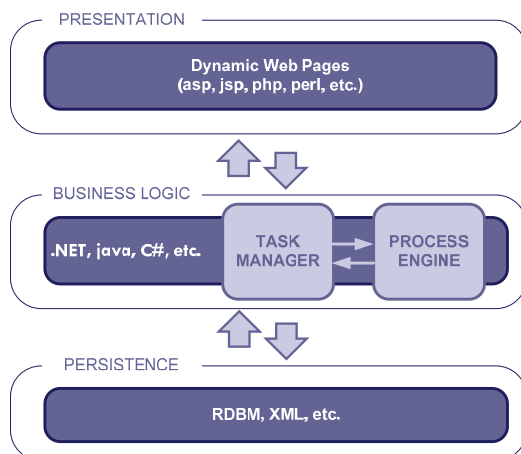


Figure 2 Extension defined at the Architectural Level.

Link:
<http://oomethod.dsic.upv.es>

Please contact:
 Vicente Pelechano
 Universidad Polit3cnica de Valencia
 SpaRCIM, Spain
 E-mail: pele@dsic.upv.es

Coordination of Emergency Communication in Safe Hands

by Annette Kik

A Belgian Hercules airplane crashed at Welschap Airport near Eindhoven, the Netherlands on 15 July 1996. Firemen extinguished the fire, unaware of the fact that over forty people were still inside. Thirty-four people did not survive. This is a sad example of human communication that went wrong in moments of stress. Would an automatic communication system have detected that essential information was missing? Within the Cybernetic Incident Management (CIM) project, CWI collaborated with universities and high-tech companies to improve communications in emergency situations.

Formerly, disaster management was all about knowing contingency plans by heart: mostly formal and legal rules. Nowadays, a more active approach is needed. During realistic simulations it often appears that communication is the bottleneck in disaster management.

tem know who is on duty to form this team and whom to call if some of them cannot be reached or are otherwise unavailable. It calls people through their preferred communication medium, such as analogue or ISDN telephone, GSM, VOIP, SMS or e-mail. It can

could be saved and some 'filtering problems' avoided if their first hand information could automatically be dispatched to the appropriate people involved. Despite all benefits, trials to experiment with it encountered some opposition in governmental organizations. Why change a safety-critical system that works well, at least most of the time?

For Almende, this was no reason to stop further product improvement. Improvements include the concept of mobile channels developed by the SEN3 research group at CWI - a change in the ideas of building modern software systems. The essentials are not about pieces of software anymore, but about programming the communication and coordination between them. Farhad Arbab, researcher at CWI and professor at Leiden University, foresees big advantages for Service-Oriented Computing, where composition of existing services can be offered as a new service.

Infrastructure

With mobile channels independent organizations can set up new businesses that do not require alterations to existing services. The crucial point is that mobile channels offer a mechanism to fully decouple software behaviour from its underlying code. Mobile channels only know dynamic connections: They determine which software module is connected to the others, and when. This becomes indispensable when each concern falls within the jurisdiction of an independent autonomous organization, as is the case in incident management.

CWI's Reo system implements communication and coordination protocols that regulate, synchronize, and combine the data streams through mobile channels. In the above example, if one volunteer



Communication is often the bottleneck in disaster management.

Reason for high-tech company Almende to investigate whether they could apply their ASK system to emergency control. This was originally designed for dynamic resource planning, communication, and distributed knowledge management.

The ASK system is based on a set of intelligent agents: autonomic pieces of software that collaborate to fulfil a certain task. For instance, if 10 volunteer firemen are needed, agents in the sys-

scale-up and escalate a situation according to the communication protocol. It searches the best solution and after the call, it asks for feedback in order to improve itself. Many test scenarios are provided. It takes care of the complete communication coordination.

The ASK system has several benefits. It is more consistent and reliable than human beings in stress situations. Workers are often closer to a disaster than a coordinator, so precious time

fireman cannot assist at a certain moment, the protocol can have the system switch to a neighbouring fire department, without the involvement or knowledge of the fireman or his fire department. The SEN3 group provides the infrastructure for these communication systems.

Future

Will the ASK system ever be used for emergencies? To test it in a more neutral and less stressful environment, the system was set up at an employment agency. Fifty freelancers received e-mails or SMS messages asking if they could work at a certain time. Without knowing that they were only talking to a computer, they all typed in the answer.

Where it took one person at the employment agency 8 hours to call and schedule 50 people, the ASK system performed the same task in less than three minutes. The CIM team hopes that these kinds of successes help to win confidence and create the opportunity to test the system within both simulated and real emergency situations. It is all about coping well with small risk situations that can have big consequences.

The CIM project started in 2003 and will end in 2007. It is being financed by SenterNovem in the Netherlands. CWI's research partners in this project are the Technische Universiteit Delft, the Vrije Universiteit Amsterdam, Almende, CMotions and Falck.

Links:

<http://www.cwi.nl/sen3/>
<http://homepages.cwi.nl/~farhad/>
<http://www.almende.com/>

Please contact:

Farhad Arbab
CWI, The Netherlands
Tel: + 31 20 592 4056
E-mail : Farhad.Arbab@cwi.nl

Super Progress in Super Computing

by Harry Rudin

Digital-computer-based analysis and simulation have long been used to solve the most difficult scientific and engineering problems. As computer power hurtles along the performance curve predicted by Moore's Law these techniques have become more and more successful – and over an ever increasing spectrum of problems. Having a glimpse at these techniques' success, highly parallel computer architectures (or super computers) have been developed to attack enormously complex problems. Super computers are even being used to improve super-computer technology. Here, as examples, a popular software package and progress in semiconductor analysis are discussed.

'Super computers' or better, high-performance computers (HPCs) use a similar technology to that in our laptop computers. However, HPC structure is massively parallel with sophisticated communication paths connecting the component processors with one another and with the processors' memory. The trick is to take advantage of the parallel processing capability without having inordinately heavy communication overhead slow the overall system down. The right architecture along with the right software can lead to gains in processing speeds of some three orders of magnitude compared with the modern laptop.

HPC applications cover a huge area. A few examples are:

- more accurate weather forecasting
- modeling and analyzing atmospheric pollutant flows
- development of new pharmaceuticals
- analysis of ocean waves breaking on sea walls
- designing better automobiles


- optimizing fire-escape routes in large buildings
- prediction of the immunological response to drugs,
- natural catastrophe damage estimation for reinsurance,
- flow of air over aircraft wings, and even
- analysis of thermonuclear fusion for power generation.

These and other applications hold enormous promise for the future. Most impressive is perhaps the last, where the hope is that, beginning in a decade's time, we will see a way to electric power without the problems of radiation from nuclear waste.

One of the techniques which is enabled by HPC is the modern application of 'Molecular Dynamics', a specially tailored form of numerical computer analysis in which the interactions of molecules and atoms with one another are simulated, based on the notions of statistical mechanics, a collection of the

relevant physical laws. To uncover physical and chemical properties of a material, a simulation of a large number of atoms or molecules, say in the order of one thousand, is required. A particularly successful software system for handling such simulations was initiated at the IBM Zurich Research Laboratory and is called CPMD for Car-Parinello Molecular Dynamics. CPMD continues to be refined there and at partner institutions. The software is distributed via a license but without cost. The CPMD code has been widely used for the analysis of physical and chemical properties of many different kinds of compounds. It has proven to be so popular that it currently has over six thousand registered users, worldwide.

Recently CPMD has been applied to the problem of analyzing a new gate dielectric for improving the performance of transistors. A most promising material is hafnium dioxide. While hafnium dioxide looked to be the ideal candidate for a new gate dielectric, there were



A partial image of a typical model of hafnium silicate used in the study.

concerns that its use might lead to unforeseen consequences in semiconductor production lines. CPMD was used to fully understand the physics behind hafnium dioxide and its behavior when used together with silicon in semiconductors. Even with the use of a high-performance computer running the efficient CPMD code, the analysis was an

enormous undertaking. Some fifty different models of the use of hafnium dioxide were analyzed. The models simulated the interactions of up to 600 atoms. Using the IBM Zurich Research Laboratory's Blue Gene/L supercomputer system with 4096 processors, five days were required to complete the analysis of each model. Thus some 250

supercomputer days were committed to this project. All of this would have been impossible, just one decade ago. But the result was well worth the effort: the developers can now sleep well.

In effect, this work completes another cycle: super computers are used to increase the speed of the processors needed to make super computers even more powerful. Quoting Alessandro Curioni, computational material scientist and supercomputing expert from IBM's Zurich lab: "So ... we are able to use supercomputers to investigate materials that will be eventually used in the next generation of supercomputers."

Links:

<http://www.cpmid.org/>
<http://www-03.ibm.com/press/us/en/pressrelease/21142.wss>
http://domino.watson.ibm.com/comm/pr.nsf/pages/news.20021212_super-computer.html

Please contact:

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

Grid Technology Makes Weather Forecasts without Boundaries a Reality

by Baudouin Raoult, Clemens-August Thole and Ute Gärtel-Zafiris

The European SIMDAT project provides the basis for a cost-effective, global information system for meteorology and the environment.

The results obtained from SIMDAT, a European research and development project, are increasingly in demand from European and international meteorological services and are likely to become acknowledged worldwide. SIMDAT Meteo is working to establish a Virtual Global Information System Centre (VGISC) for the national meteorological services of France, Germany and the UK. This will be based on Grid technology and will be used within the World Meteorological Organization Information System (WIS) to provide cost-effective and user-friendly services. VGISC offers a unique meteorological database integrating a variety of data and providing secure, reliable and convenient access via the Internet. It is targeted towards

operational services and research in the domains of meteorology, hydrology and the environment.

A Grid is defined as a software system that provides uniform, location-independent and secure access to geographically and organizationally scattered and varied resources (such as databases, analysis services and computational power). The objective of the SIMDAT project – which has 11 million euros of funding support from the European Commission and is coordinated by the Fraunhofer Institute for Algorithms and Scientific Computing (SCAI) – is to accelerate the uptake of Grid technology not only in meteorology, but also in areas such as the automotive, aerospace and pharmaceutical industries.

The VGISC software, developed by the SIMDAT Meteo project partners and led by the European Centre for Medium-Range Weather Forecasts (ECMWF), will offer meteorological communities worldwide immediate, secure and convenient access to various data and analysis services, as well as a user-friendly platform for storage of meteorological data. VGISC will thus enable the fast exchange of data for numerical weather forecasts, disaster management and research, while remaining independent of national frontiers and beyond organizational boundaries.

Weather does not recognise borders. To study historical data regarding the German low mountain ranges, for example, one only needs to analyse data from the

Deutscher Wetterdienst (DWD). However, for a similar study in the German Alps, one must obtain data from neighbouring states, since a centralized catalogue for meteorological data gathered from all the different countries is still lacking.

It is in this type of situation that VGISC, as part of WIS, can offer new solutions. The new system is a cost-effective, reliable and highly flexible 'one-stop-shop' that combines available meteorological data into one system. This catalogue will replace the current multiplicity of systems, which frequently result in incompatibilities, inefficiencies and the duplication of efforts.

The infrastructure of this new system will be based on a mesh network of peers and meteorological databases. Messages are interchanged using algorithms based on mobile telephony technology and metadata synchronization on a journalized file system. The Grid technology is based on Open Grid Services Architecture Data Access and Integration (OGSA-DAI™), which is founded on Web service and Web technology concepts. In addition, standard protocols such as Open Archive Initiative (OAI) are used to synchronize and integrate existing archives and databases as well as to extend interoperability. Furthermore, VGISC will act as a testbed for the ISO 19115 metadata standard by handling complex data in real time.

The SIMDAT project is Europe's contribution to the infrastructure technology of the emerging WIS, as the World Meteorological Organisation (WMO) modernizes and enhances its long-standing Global Telecommunications System (GTS), an international network for exchanging mainly meteorological data and warnings in real time. In addition, the new system will provide access to all environmental communities worldwide, whereas the GTS only gives access to the existing national weather services of the member states.

The opportunities for the new VGISC technology are excellent, as VGISC is not only of interest within Europe: the national meteorological services of Australia, China, Japan, Korea and the Russian Federation's National Oceanographic Centre have already deployed the SIMDAT software and are collaborating actively with the European part-

SIMDAT **CECMWF** Home Search Directory Requests Keyword search

Metadata Home > Metadata >

[Retrieve Data](#) [Show XML](#)

Title: *Freely available data from ECMWF 15 Years Re-Analysis*

Abstract: This dataset contains daily fields, monthly means, monthly daily means, invariant fields and monthly invariant fields from ECMWF 15 Years Re-Analysis, on a 2.5 degree grid. Pressure level parameters: Geopotential, Temperature, U velocity and V velocity. Single level parameters: 10 metre U wind component, 10 metre V wind component 2 metre dewpoint temperature and 2 metre temperature and Total cloud cover.

Period: 1979-01-01 to 1993-12-31

Bounding Box: 90°N 180°W 90°S 180°E

Categories: [NWP Outputs](#) > [ECMWF](#) > [15 years reanalysis](#) > [Centre](#) > [ECMWF](#) > [15 years reanalysis](#) >

ECMWF Analysis VT: Sunday 1 September 1957 12UTC Surface: 2 metre temperature

Discovery and access to meteorological data through an internet portal

ners. The software deployment is followed by an increasing number of meteorological centres, and new meteorological datasets from Asia, Australia, Europe and the United States are steadily being added to the portal.

VGISC software has been developed by the ECMWF, DWD, Météo France, the UK Met Office, EUMETSAT and other SIMDAT partners. The SIMDAT consortium comprises 25 European enterprises and research centres from various disciplines. The project is coordinated by the Fraunhofer Institute SCAI in Sankt Augustin, Germany.

The Fraunhofer Institute for Algorithms and Scientific Computing engages in computer simulations for product and process development and is a strong partner in industry. SCAI designs and optimizes industrial applications and makes calculations on high-performance computers. Fraunhofer Gesellschaft undertakes applied research of direct use to private and public enterprise.

The European Centre for Medium-Range Weather Forecasts (ECMWF) is

an international, inter-governmental organization, supported by 28 European states. The objectives of the Centre include the provision of medium-range forecasts to the meteorological offices of its member states and cooperating states, maintaining a data archive, and providing assistance in advanced education and support to the World Meteorological Organization in implementing its programmes.

Links:

<http://www.simdat.eu>
<http://www.ecmwf.int>

Please contact:

Clemens-August Thole
and Ute Gärtel-Zafirris
 Fraunhofer SCAI, Germany
 Tel: +49 2241 14 2178
 E-mail: simdat@scai.fraunhofer.de

Baudouin Raoult
 European Centre for Medium-Range
 Weather Forecasts, UK
 Tel: +44 118 949 9404
 E-mail: baudouin.raoult@ecmwf.int

Meeting the Multicore Challenge: Dynamic Dependence Analysis for Parallelization

by Karl-Filip Faxén

Parallel programming is no longer optional. As improvements in processor clock frequencies have levelled out, the only way to increase performance is to place several processor cores on each chip. In contrast to increases in clock frequency, multiple cores do not automatically help existing programs, a problem addressed by a new multicore initiative from SICS.

While offering improved performance, multicore processors pose new demands on software; existing applications will in general not see any improvements at all. This is because each core on a multicore chip is no more powerful than a traditional processor; the performance improvements come from using more than one core at the same time. For this to happen, each core must have its own program. If one wishes to run a single application faster, then that application must be divided into sub-programs, or threads, that cooperate to deliver the desired functionality.

Traditionally, most programs consist of only one thread. This is no accident; it is easier to think about a program that does one thing at a time than at a program that does several things at a time. So not only do we need to rewrite a lot of code; we need to make it parallel, which is currently very difficult. Thus we need better ways both to write new parallel programs and to make our legacy code threaded.

The main obstacle in this partitioning process is dependencies. If one part B of the program has the output of another part A as input, then A must run before B. Other dependencies arise if two parts of the program use the same memory area for temporary storage: this means they cannot run in parallel. Other kinds of dependencies arise through I/O and when the execution result of one part of a program determines whether another part should be executed at all (think of conditionals).

SICS, as an industrial research institute, is starting a proactive program for cooperation with software-intensive Swedish industry in this area. Our work aims to help developers find the potential for parallelism in programs, in particular by providing efficient tool support. We are currently working on

Embla, a dynamic dependence analyser capable of finding opportunities for manual parallelization of sequential applications.

We are interested in the following methodology for constructing parallel programs. Start from a sequential program (old or new), identify independent parts of that program using Embla, and rewrite the program to obtain parallel execution of the independent parts. The

sible for selecting program inputs that generate representative program executions with good coverage.

If all dependencies are found, and the parallelizing transformations made do not violate these, the parallel program is guaranteed to behave like the sequential version. Conversely, if the parallel program behaves differently, the reason must be that a dependence was missed. The sequential program can then be run

```

1:      #include <stdlib.h>
2:      #include <stdio.h>
3:
4:      static void inc(int *p)
5:      ...{
6:          *p=*p+1;
7:      ...}
8:
9:      int main(int argc, char **argv)
10:     ...{
11:         int *q=NULL,n=0;
12:
13:         q = (int*) malloc( sizeof(int) );
14:         inc(q);
15:         inc(&n);
16:         inc(q);
17:         printf( "%d\n", *q+n );
18:         q = (int*) malloc( sizeof(int) );
19:         return q==NULL;
20:     ...}

```

The output of Embla for a small C program. Arrows represent dependencies between source lines; line 15 is independent of lines 13-16 and can be run in parallel.

figure shows the output of Embla for a trivial program with dependencies shown as arrows. Note that while a dependence exists between line 14 and line 16 since the two calls increment the same variable, both of these lines are independent of line 15.

Embla observes the actual data dependencies that occur during program execution, decides which program parts are affected by the dependence and presents the results. Since Embla observes particular runs with particular inputs, it might happen that when the program is run with other inputs, other dependencies occur. The user of Embla is respon-

under Embla again with the offending input, revealing the missed dependence. Thus, we need never debug the parallel program itself; all debugging is made in the sequential context.

We also plan activities ranging from research on parallelization and execution environments for multicore processors to more applied projects in close cooperation with industry.

Please contact:
Karl-Filip Faxén
SICS, Sweden
Tel: +46 8 633 1610
E-mail: kff@sics.se

Multimodal Multi-Level Fusion using Contextual Information

by Olga Vybornova, Monica Gemo and Benoit Macq

A new method of multi-modal, multi-level fusion that integrates contextual information obtained from spoken input and visual scene analysis is being developed at Université Catholique de Louvain in the frame of the EU-funded SIMILAR Network of Excellence. An example for application is the 'intelligent diary', which will assist elderly people living alone to perform their daily activities, prolong their safety, security and personal autonomy, and support social cohesion

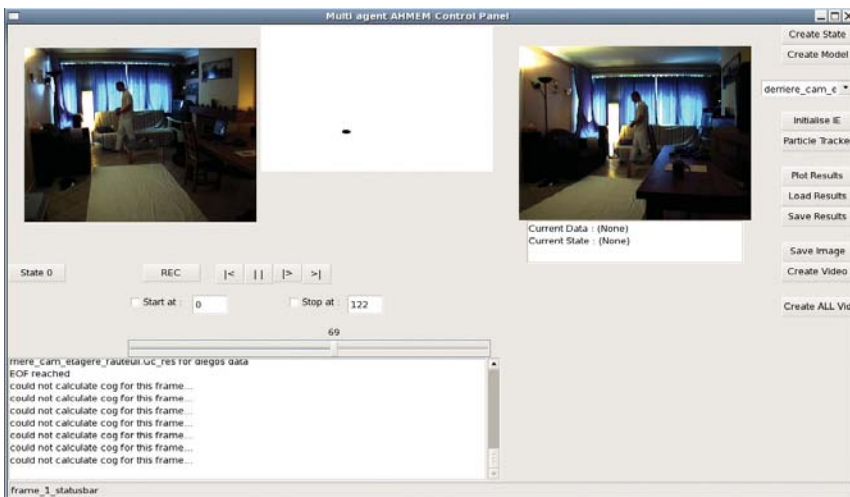
A system whose objective is to interact naturally with humans must be able to interpret human behaviour. This means extracting information arriving simultaneously from different communication modalities and combining them into

the same multi-modal interface (ie an interface using at least two different modalities for input and/or output), it is necessary to extract relevant features from signal representations and thereafter to proceed to high-level or

The SIMILAR network of excellence provides algorithms that integrate common meaning representations derived from speech, gesture and other modalities into a combined final interpretation. The higher-level fusion operation requires a framework of common meaning representations for all modalities, and a well-defined operation to combine partial meanings arriving from different signals. In our case, the resulting fused semantic representation should contain consistent information about the user's activity, speech, localization, physical and emotional state, and so on.

Everything said or done is meaningful only in its particular context. To accomplish the task of semantic fusion we are taking into account information obtained from at least three contexts. The first is domain context: personalized prior knowledge of the domain such as predefined action patterns, adaptive user profiles, situation modeling, and a priori developed and dynamically updated ontologies for a particular person that define subjects, objects, activities and relations between them. The second is linguistic context, which is derived from the semantic analysis of natural language. Finally, visual context is important: capturing the user's gesture or action in the observation scene as well as eye-gaze tracking to identify salient objects of the activity.

To derive contextual information from spoken input we extract natural language semantic representations and map them onto the restricted domain ontology. This information is then processed together with visual scene input for multimodal reference resolution. The ontology allows the sharing of contextual information within the domain and serves as a meta-model for Bayesian networks that are used to analyse and combine the modalities of



Multimodal semantic integration of speech and visual streams.

one or more unified and coherent representations of the user's intention. Our context-aware user-centered application should accept spontaneous multi-modal input – speech, 3D gestures (pointing, iconic, possibly metaphoric) and physical action; it should react to events, identify the user's preferences, recognize his/her intentions, possibly predict the user's behaviour and generate the system's own response. In our case, namely the intelligent diary, we have a restricted domain of application, but must deal with unrestricted natural human behaviour – spontaneous spoken input and gesture.

Multimodal fusion is a central question that must be solved in order to provide users with more advanced and natural interaction support. To understand and formalize the coordination and cooperation between modalities involved in

'semantic' fusion. High-level fusion of modalities involves merging semantic content obtained from multiple streams to build a joint interpretation of the multi-modal behaviour of users.

Signal-level and semantic-level processing are deemed to be tightly inter-related, since signals are seen as bearers of meaning. In practical applications it is convenient to distinguish between the two levels of abstraction. In many situations there is a tradeoff, and higher efficiency can be achieved with a better balance between low-level (signal) fusion and high-level (semantic) fusion. To tackle this problem we are complementing our initial approach on optimal feature selection, developed in the information theory framework for multi-modal signals, with a knowledge-based approach to high-level fusion.

interest. With the help of probabilistic weighting of multi-modal data streams we obtain robust contextual fusion. We are thus able to recognize the user's intentions, to predict behaviour, to provide reliable interpretations and to reason about the cognitive status of the person. Figure 1 illustrates an example of multi-modal semantic integration of visual and speech input streams. The system interprets the user's behaviour (action) after analysing contextual information about the user's location (visual context) and intended goal, as derived from the spoken utterance.

Based on our experimental tests, we suggest that in order to make the multi-modal semantic integration more effi-

cient and practical, special attention should be paid to the stages preceding the final fusion stage. The strategy here is to use visual context simultaneously with speech recognition. This helps to make speech recognition more accurate and thereby obtain a text of proper quality.

Our current research is devoted to the implementation of multi-level cross-modal fusion, which looks promising from the point of view of resolving reference ambiguity before the final fusion. It is cross-modal fusion that aids in speech recognition for the elderly, where problems are caused by age-related decline of language production ability (eg difficulties in retrieving appropriate (familiar) words, or tip-of-

the-tongue states when a person produces one or more incorrect sounds in a word). This is possible because information from other modalities refines the language analysis at early stages of recognition. Methods for achieving robust and effective semantic integration between the linguistic and visual context of the interaction will be further explored.

Link:

<http://www.similar.cc>

Please contact:

Olga Vybornova, Universite Catholique de Louvain, Belgium

Tel: +32 10 47 81 23

E-mail: vybornova@tele.ucl.ac.be

The SAPIR Project: Executing A/V Complex Queries in Peer-to-Peer Systems

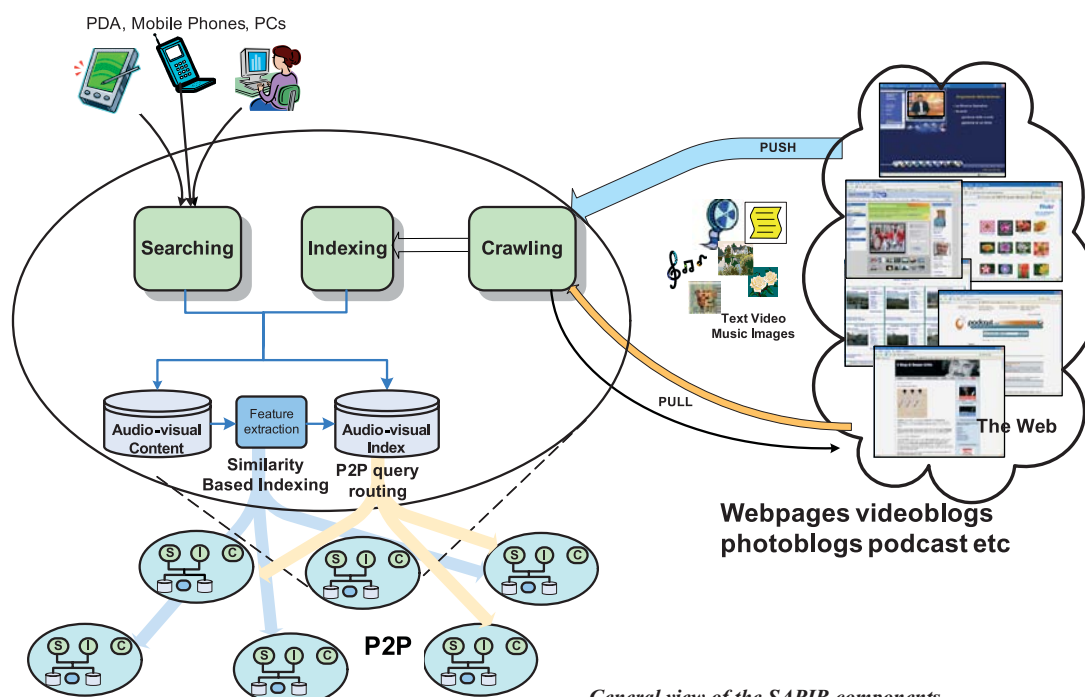
by Claudio Gennaro, Raffaele Perego and Fausto Rabitti

Searching for non-text data (eg, images) is mostly done by means of metadata annotations or by extracting the text close to the data. However, supporting real content-based audio-visual search, based on similarity search on features, is significantly more expensive than searching for text. Moreover, the search exhibits linear scalability with respect to the data set size. The European project SAPIR is currently addressing this problem.

A large component of Web content consists of non-text data, such as images, music, animations, and videos. Current search engines index Web documents by their textual content. For instance, Web

tools for performing image searching (such the ones provided by Google, Yahoo!, or MSN Live Search) simply index the text near the image and the ALT attribute of the IMG tag, used to

provide a description of an image. One reason is that search at the level of features (such as color histograms or shapes) exhibits linear scalability with respect to the data search size, which is



General view of the SAPIR components.

not acceptable for the expected dimension of the problem. The motivation is that for this kind of data the appropriate search methods are based on the similarity paradigm, which typically exploits range and nearest neighbor queries. These queries are computationally more intensive than exact match, since conventional inverted indexes used for text are not suitable for such data.

Peer-to-peer (P2P) systems are currently considered a promising way to address the problems of scalability, and several scalable and distributed search structures have been proposed covering even the most generic cases of metric space searching. A common characteristic of all these existing approaches is the autonomy of the peers with no need for central coordination or flooding strategies. Since there are no bottlenecks, the structures are scalable and high performance is achieved through parallel query execution on individual peers.

The European project SAPIR (Search on Audio-visual content using Peer-to-peer Information Retrieval) aims at

developing a P2P architecture able to provide a scalable indexing structure that can be used for dealing with complex queries, ie, queries that involve more than one feature, such as: find all images in database similar to the query image with respect to the color and the shape.. SAPIR will support multimedia content for uploading/pushing, and for searching/retrieving from a variety of devices, including mobile phones, PDAs, and PCs. In general, there will be two different types of crawling methods: pull and push. The former is the more traditional: crawlers are responsible for locating, browsing, and gathering the web resources. In the latter case, the resources upload their content directly to the crawler (see Figure 1). User context such as GPS position, query history, and social networking (for groups of users with similar interests) will be used to increase result precision. Caching techniques will be developed to increase system performance. The indexing subsystem extracts the features from the data and, based on the indexing policy, indexes them in its local "audio-visual index" or distributes them over the P2P network to remote indices. Each peer of the P2P network

provides a searching service, which either executes the query locally on its audio-visual index or submits the query over the P2P network to remote peers.

Complex queries are of fundamental importance in the application area of the project. For instance: find all images in database similar to the query image with respect to the color and the shape. In this situation, it is necessary to use an incremental nearest neighbor algorithm, since we do not know how many neighbors must be retrieved before one is found that satisfies the conditions. In the context of this project, ISTI-CNR is developing a first distributed incremental nearest neighbor algorithm for P2P-based systems. Our solution, based on a generalization of the priority queue algorithm proposed for hierarchical centralized structures, is optimal and not tied to a specific P2P architecture.

Link:
<http://www.sapir.eu/>

Please contact:
Claudio Gennaro
ISTI-CNR, Italy
E-mail: claudio.gennaro@isti.cnr.it

Getting your Data Back by Giving it Away

by Jacco van Ossenbruggen

CWI teams up with an Amsterdam-based consortium to develop web services on open linked data in the domain of cultural heritage. The approach is gaining momentum, but remains challenging from both the researchers' and museum's perspective.

The Dutch MultimediaN e-culture project is applying the idea of open linked data to the traditionally closed and isolated worlds of cultural heritage collections. Under the hood, new Semantic Web technology is developed to realize new search services and smarter Web interfaces that provide access to multiple, information rich, museum collections.

All features developed within the project are, however, judged by the way they succeed in conveying the main principle underlying the project: by providing the user meaningful relationships across collections, all individual data collections grow in value. In the same way explicit hyperlinks add value to all documents being linked to and from,

explicit relationships on the Semantic Web add value by providing context to previously isolated data.

The project builds on the fact that Amsterdam happens to be the home of three research institutes with world class Semantic Web expertise: Vrije Universiteit, University of Amsterdam and CWI. The project team is completed by representatives from two key Dutch cultural heritage institutes: ICN and DEN. The project closely cooperates with a growing list of museums. It is financed by the Dutch natural gas reserves through the national government's BSIK program, and started in 2004. From the beginning, all partners have been cooperating closely, in a way that is rarely seen in computer science research projects.

Comparable projects typically spend the last few months before the end of the project to build a proof-of-concept prototype that demonstrates that the innovations developed by the individual project partners actually work together. This project takes the very opposite route, and released the first version of an integrated prototype even before the first PhD student started to work on the project late 2005. As a result, all senior and junior researchers involved in the project use – and contribute to – the same experimentation platform from the very beginning.

The platform is based on the open source SWI-Prolog package, and all generic Semantic Web technology developed within the project is released

as part of the standard distribution. Conversion software has been developed to convert museums' collection databases, thesauri and other sources of domain knowledge to RDF, and to create meaningful links between the sources. For CWI, the main research challenges are the design and evaluation of new search functionality that is made possible by the linked data, and the design and evaluation of the associated web interfaces.

The project's approach payed off immediately: eight months later, in August 2006, the second release of the platform was submitted to the International Semantic Web Challenge, a submission that turned out to be a winning one during the International Semantic Web Conference in November 2006. To project's impact goes, however, well beyond the computer science research world. Less than five months after the ISWC award ceremony, the project was presented before the international cultural heritage community, which was gathered in San Francisco in April 2007 for the Museums and the Web conference. There it turned out that the project had made the right decision to base its strategy on open and linked data. Museums and archives all over the world start to realize the serious limitations of vendor lock-in and closed proprietary solutions. While we researchers think about



Project leader Guus Schreiber (right) demonstrating the project to Tim Berners-Lee, Director of the World Wide Web Consortium, at ISWC 2006.

ways to provide a wide audience better access to the museum collections over the public Web, on the other side of the firewall, many museums fight to get access to their very own data, as it is locked inside proprietary software without public APIs.

In this context, it is clear that a future in which data is open may scare off many in the museum world. It takes time to get used to the idea that your own website is no longer the only way to access 'your' data. But without a doubt, sooner or later both the museum itself, unknown users and third parties will develop a wide range of new web services, mash-ups, widgets, social tagging applications, and much more based on the museum's data. This is simply

because what the curators call 'their data represents 'our' heritage, and there are just too many of us users interested in these rich information resources. Providing a wider public access to data in a commonly agreed upon, open and linkable format is definitely the way to go. It may even be the only way to get the large amounts of valuable data back that is now locked up inside proprietary formats.

Link:

<http://e-culture.multimediam.nl/>

Please contact:

Jacco van Ossenbruggen

CWI, The Netherlands

Tel: +31 20 592 4141

E-mail: Jacco.van.Ossenbruggen@cwi.nl

OPTIHPER: A Computer-Based Decision Support System for Employee Timetabling Problems

by Antonio Lova, Pilar Tormos and Federico Barber

A high degree of job satisfaction is an essential factor in the success of a company. A company's objectives, labour conditions and worker preferences must be optimized in a complex space of solutions. OPTIHPER (Optimización de Horarios de Personal - Employee Timetabling Optimization) is a computer-aided system able to efficiently assign employees to tasks while verifying a wide set of constraints and optimizing organizational objectives and employee preferences. The system is currently in use by leading commercial companies with very good results

Employee Timetabling Problems (ETPs) arise in any organization with a set of tasks that must be assigned to a set of employees, each with their own qualifications, constraints and preferences. This problem arises in many institutions (hospitals, supermarkets etc) where the assignment of employees is usually performed manually. However, in companies with a large number of workers and

working centres, the use of computer-based tools is essential for making organizational objectives compatible with worker preferences.

ETPs can be formulated as a constraint network, where the main entities that define the problem are tasks, workers and timetables. Tasks can be sporadic or can be repeated along one shift.

They can have either a fixed duration or require a certain amount of work (man-hours) to complete. Workers may have different degrees of qualification for each task and can be assigned to different timetables.

Timetables are defined by the start/finish time and the shift (usually morning, afternoon, evening and night).

tion is found in the local search phase. OPTIHPER has two execution modes:

- 'Standard mode' computes optimized task assignments considering all available workers,
- 'Opti-Staff mode' does so for a limited number of workers, thereby allowing us to determine the optimum size/qualification of the staff.

Architecture of OPTIHPER

OPTIHPER is a multi-platform system implemented in standard ANSI C. It takes the required data either from standard databases or input files (see Figure 1). After its execution, OPTIHPER generates several statistical and graphical reports (see Figures 2 and 3).

Conclusions

OPTIHPER integrates AI/OR techniques able to cope with complex ETPs arising in the real world. It can obtain optimized assignments for when the entire staff is available or when the number of workers is fixed below the total available. Means technological transference agreements, customized versions of this system, are in operation with excellent results in leading European distribution companies. Standard scenarios involve the assignment of over one hundred workers, fifty types of task and sixty labour timetables. Assuming a planning horizon of four weeks, this implies personnel assignments of more than 15,000 hours. The

optimized solution is obtained in a few seconds. The performance of the system, along with its flexibility and efficiency, make OPTIHPER an essential tool for staff allocation in many contexts, allowing companies to increase their competitiveness and to optimize their workers' skills and preferences.

Link:

<http://www.dsic.upv.es/users/ia/gps/optihper/>

Please contact:

Federico Barber
Polytechnic University of Valencia
Tel: +34 96 3879357
E-mail: fbarber@dsic.upv.es

First EchoGRID-EUChinaGRID International Conference

The First International Conference on European Union and Chinese Grid Experiences, held in Beijing on 24-25 April, provided an outstanding platform for European and Chinese state-of-the-art Grid technologies and projects, showcasing results, new application areas and future advancements. The conference attracted some 150 experts and professionals from large enterprises, SMEs, academic and research organisations, government and public administration, international standardisation bodies, scientific and business associations from Europe, China and beyond. The event was organised by EchoGRID and EUChinaGRID, two projects funded by the European Commission, and hosted by the Institute of Computing Technology, Chinese Academy of Science.

The first day of the conference set the scene for ensuing discussions on themes of mutual interest and common challenges.

Dr Xiaohan Liao, Deputy Director, High & New Technology Department, of the Chinese Ministry of Science and Technology (MOST) addressed over 150 attendants from research, business, government and scientific organisations. Alison Birkett, European Delegation to China, was delivering a welcome address highlighting the delegation's mission to extend and intensify the dialogue and co-operation in all the EU's areas of competence, particularly in ICT research and the EU-China Information Society Project.

Keynote addresses were delivered by Chuncheng Wang, Division of Information, MOST, offering insight into Grid Research in China's High-Tech R&D Programme. The participants then heard from four key industrial actors: Wenbo Mao (HP Lab China); Xintai

Wang (eStarCom Inc), Wayne Wang (Intervision Software), and Donpu Fu (TongTech).

Thematic Session 1 'Enterprise Challenges with Grids' was led by key business experts. Professor Jun-Seok Hwang, Seoul National University and head of several industrial associations, reported on the major challenges for Open Grid Services. Pawel Plaszczak, president of GridwiseTech (Poland), presented a series of fascinating Grid Technology case studies highlighting benefits for businesses. Dr Wei Zhou, CNIC, CAS, focused his presentation on the new context for innovative interdisciplinary research offering added value to R&D activities.

The afternoon session of the first day opened with a key theme 'Interoperability'. Yongjian Wang, Beihang University, evaluated the challenging and multi-faceted issues related to interoperability, such as job management, data management, workflow and security.

The mission of EUChinaGRID is to implement interoperability between CHGrid GOS and EGEE gLite. This talk presented the gateway-based solution proposed by EUChinaGRID aimed at implementing batch level job interoperability and providing a prototype to test it.

The presentation by Jianxin Li, also from Beihang University, revolved around the security architecture for CROWN (China Research & Development Environment over wide-area Network). The talk discussed the planned extensible framework enabling distributed access control & dynamic trust establishment between service providers and consumers in a Grid environment. CROWN Grid research into policy negotiation & virtual organisation security management was also discussed.

Dr. Gang Chen, IHEP, offered insight into LCG, a Grid project aiming to provide HEP computing infrastructure for Large Hadron Collider (LHD) experi-

ments at CERN (Switzerland). Dr. Chen gave an overview of the architecture and design of the LHC computing Grid, and reported on experiences in deploying LCG in China.

The afternoon concluded with six demos:

- Maciej Malawski, Jagiellonian University (Poland): Biological applications in EUChinaGRID
- Zhao Yongwang, Beihang University (China): Seismic Data Processing
- Boqun Cheng, ICT (China): the VEGA PROJECT
- the CROWN Team (China): CROWN
- Clement Mathieu, Denis Caromel, and Yu Feng, INRIA (France): the ProActive Java GRID Middleware Library.

The second day started with a session dedicated to new programming paradigms. It included a talk by Professor Zhiwei Xu, Institute of Computing Technology, CAS, entitled 'Network Centric Operating Systems'. Professor Huamin WANG, NUDT (China), gave a presentation on Internet Virtual Computing Environment (IVCE), while Denis Caromel, INRIA (France) discussed Grid component models and active objects.

Andrea Manieri, Engineering (Italy) presented the NESSI vision on research challenges. The European Technology Platform, NESSI (Networking Software and Service Initiative) aims to create synergies and help pave the way for the future Service-based Economy for Europe. Enterprises and research organisations have come on board this project and its working groups to discuss future research challenges. NESSIGrid will

foster collaboration in the Grid Research Area and support NESSI in the definition of a strategic research agenda for the context of Grid and Service-Oriented Infrastructure.

The talk by Zsolt Nemeth, SZTAKI (Hungary), looked at a highly abstract coordination model for distributed workflow enactment where decentralised control, autonomy, adaptation to high dynamics and partial lack of information are primary concerns. The talk illustrated how this model provides a framework where complex workflow patterns and advanced issues of enactment can be modelled.

The thematic session on management in Grids highlighted the increasing need for management mechanism enabling Grid computing to satisfy diverse application requirements. The presentation by Professor Zhongzhi Luan, Beihang University, explained why management should be an essential attribute of Grid on account of the large-scale Grid environment and rapid extension in terms of both resources and networking. Management mechanisms need to be flexible and automated. Professor Luan discussed the state of the art, key issues and useful technologies and related research activities. The talk by ETICS project representative, Andrea Manieri, revolved around a quality certification model for Grid research projects, showcasing the ETICS feasibility study. ETICS is aimed at developing a Grid-based infrastructure for building and testing distributed software. ETICS tools take care of the whole integration-build-testing-packaging phase crucial for many research projects.

Chinese national Grid initiatives were the focus of the session dedicated to ongoing research versus enterprise challenges. The three initiatives in the spotlight were ChinaGrid, CROWN and CNGrid.

A roundtable on future collaborative scenarios discussing and evaluating key areas for future co-operation for innovation, closed the conference.

During the conference, INRIA provided a two-day free ProActive tutorial. ProActive has been successfully used as a strong vehicle for interoperability in the first three Grid Plugtests events, and it is also the current reference implementation of the Grid Component Model (GCM). Moreover, as an open source middleware part of the OW2 consortium (<http://www.ow2.org>), ProActive is being used by several industrial companies in production applications.

The conference was co-organised by ERCIM, coordinator of the EchoGRID project. A second EchoGRID workshop on European-Chinese collaboration in Grid research and technologies will be held in conjunction with the Grids@Work 2007 Joint European Union - China GRID Days in Beijing, on 28 October to 2 November 2007 (see announcement on page 70).

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

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

ERCIM-Sponsored Events

ERCIM sponsors up to ten events per year (conferences, workshops and summer schools). The funding for all types of events is in the order of 2000 Euro. Upcoming events sponsored by ERCIM include:

- *ECOOP'07 - European Conference on Object Oriented Programming, Berlin, Germany, 30 July - 3 August 2007*
- *ESSLLI 2007 - 19th European Summer School in Logic, Language and Information, Dublin, Ireland, 6-17 August, 2007*

- *ESEC/FSE Joint meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, Dubrovnik, Croatia, 3-7 September 2007*
- *4th IEEE International Symposium on Wireless Communication Systems 2007, Trondheim, Norway, 17-19 October 2007*
- *SOFSEM 2008 - 34th International Conference on Current Trends in Theory and Practice of Computer Science, Nový Smokovec, High Tatras, Slovakia, 19-25 January 2008*

More information: <http://www.ercim.org/activity/f-events.html>

CALL FOR PARTICIPATION



International Brokerage Event on R&D & Technology Transfer in Human-Centric ICT

Winterthur/Zurich, Switzerland, 16 November 2007

This networking event, sponsored by the ERCIM member SARIT (www.sarit.ch), will focus on the two ICT challenges forming the core of the FP7-ICT Call 3 to be launched around Christmas 2007:

- Challenge 2: cognitive, interactive, robotics and context-aware systems
- Challenge 4: knowledge & content interactivity, including digital libraries & technology-enhanced learning

In addition, thanks to the cooperation of the Innovation Relay Centres, this event will as well provide you with opportunities to offer or find new technologies related to the above challenges. Finally, the programme will also provide dissemination of recent results. Seize opportunities provided to make this event successful for you:

- keynote speakers from the European Commission and from FP6 projects
- short pitch presentations on project ideas for ICT Call 3
- Prearranged face-2-face 30 min matchmaking meetings for Technology Transfer
- demos and posters to animate the informal networking throughout the Day.

Deadline

Deadline for registrations for your networking contributions: 8 October 2007.

More information:

<http://www.ictsummit.eu>, ictsummit@euresearch.ch

CALL FOR PARTICIPATION

ECSA 2007 - First European Conference on Software Architecture

Aranjuez (Madrid), Spain, 24-26 September 2007

The role of software architecture, in the engineering of software-intensive applications, has become more and more important and widespread. Component-based and service-

oriented architectures are key aspects to the design, development and evolution of all software systems.

The European Conference on Software Architecture is the premier European conference dedicated to the field of software architecture, covering every architectural aspect of software and service engineering. It is the follow-up of a successful series of European workshops on software architecture held in United Kingdom in 2004 (Springer LNCS 3047), Italy in 2005 (Springer LNCS 3527), and France in 2006 (Springer LNCS 4344). Built on their success, it has evolved into a full-fledged series of European conferences whose first edition is ECSA 2007 in Madrid, Spain.

ECSA 2007 will bring together researchers and practitioners from academia and industry, to share novel ideas on the foundations, languages, models, techniques, tools, and applications of software architecture technology. It will include keynote papers, research, experience, and challenge papers, posters, and panels.

As key speakers the ECSA 2007 will include Prof. David Garlan, Carnegie Mellon University, United States; Prof. Ron Morrison, University of St. Andrews, UK; and Prof. Mike Papazoglou, University of Tilburg, The Netherlands.

The conference is organised by the Kybele Research Group, Rey Juan Carlos University, Madrid in cooperation with the World Wide Web Consortium.

More information:

<http://contraintes.inria.fr/CSCLP07/>

CALL FOR PARTICIPATION

PACT 07 - Parallel Architectures and Compilation Techniques

Brasov, Romania, 15-19 September 2007

The purpose of PACT is to bring together researchers from architecture, compilers, applications and languages to present and discuss innovative research of common interest.

Currently confirmed keynote speakers include:

- Yale Patt, University of Texas, Austin
- Bjarne Stroustrup, Texas A&M University.

Workshops

The conference hosts the following workshops

- GREPS: GCC for Research in Embedded and Parallel Systems
- Operating Systems for Heterogeneous Multicore Architectures (OSHMA)
- MEDEA: MEMory performance: DEaling with Applications, Systems and Architecture
- Massively Parallel Algorithms for Radiation Transport and Radiation Hydrodynamics.

Tutorials

- Microarchitecture: Concepts, Tradeoffs, the Future
- Title TBD
- CellSim: a Modular Simulator for Heterogeneous Chip Multiprocessors
- Transactional Programming in a multi-core Environment.

The organisers have received 175 paper submissions. The accepted papers presented at the conference will be available on the conference web site.

PACT '07 is dedicated to the memory of Irina Athanasiu. Irina was a computer science professor at the Politehnica University, Bucharest. She taught, mentored and advised generations of Romanian computer scientists who have achieved prominent positions in academia and industry throughout the world. She dedicated her life to her students and our profession.

More information:

<http://pact07.cs.tamu.edu/>

CALL FOR PAPERS

DECOS/ERCIM Workshop on Dependable Embedded Systems

In conjunction with SAFECOMP 2007, Nuremberg, Germany, 18 September 2007

“Dependable Embedded Systems – Challenges, Impact, Solutions, Examples, Professional and Academic Education and Training” is the topic of the 2007 workshop. Topics include:

- design concepts and architectures for dependable networked embedded systems
- methods, means and techniques to tolerate, remove, to prevent and forecast faults in dependable networked embedded systems
- functional Safety and Security Standards, validation and certification of dependable networked embedded systems
- tools and tool chains to facilitate design, development, operation and maintenance effectively and efficiently of dependable networked embedded systems in industrial context
- education and Training issues, means and methods to cope with the growing demand for people being aware and professionals in a holistic way of thinking

Deadline for paper contributions: 6 August 2007

Papers will be peer reviewed and proceedings be published by ERCIM after the workshop, including results of discussions as well.

More information:

<http://www11.informatik.uni-erlangen.de/safecomp2007/>

<http://its.arcs.ac.at/ercim/download/>

[CfP_SAFECOMP_2007_DECOS_ERCIM_Workshop.pdf](http://www11.informatik.uni-erlangen.de/safecomp2007/CfP_SAFECOMP_2007_DECOS_ERCIM_Workshop.pdf)

CALL FOR PAPERS

STV'07 - Fifth International Workshop on System Testing and Validation

Paris, 5 December 2007

The need to count on trusted systems is growing day by day. Together with traditional domains such as medical devices, automotive, railway, aeronautical, space and telecommunications, new applications and services are coming up everyday. Terms such as ubiquitous, pervasive, or autonomic computing, products under the general umbrella of smart devices or the use of large wireless sensors networks indicate a clear trend in the increase of application complexity and dependency. The dependency of daily life on computers and computer based systems is growing up at a high rate, and validation is gaining importance at the same rate. To perform a proper validation process encounters a number of problems both from a technical and from a managerial point of view, also considering that the complexity of the underlying software is growing as well. As a consequence inputs from research contributions and experience from industry are required in order to enable innovative, and often more rigorous, approaches.

The System Testing and Validation Workshop is a series of events initiated in the year 2002 and seeks to provide answers to the many open issues related to validation.

The objective of the workshop is to bring together industry and academy to debate on different approaches and methodologies to system validation. Issues to be considered are:

- innovative approaches to validate 'fit for use' of different types of software intensive systems with different requirements: functional and non-functional (reliability, safety, maintainability, etc.);
- automation and tool support: validation for systems of formerly-validated components;
- validation in the context of different lifecycle process models: eg conventional and agile;
- validation versus 'qualification' and certification.

Both methodological issues and managerial (cost, organization required) are welcome: contributions that put special emphasis to the usage and suitability of advanced languages for tests and test methodologies (eg TTCN-3, UML2 testing profile), practical aspects of validation (though their foundation will be equally considered).

Submission deadline: 27 August 2007

The workshop is part of the 20th International Conference "Software & Systems Engineering and their Applications" (ICSSEA 2007).

More information:

<https://syst.eui.upm.es/conferences/stv07/>

<http://www.icssea2007.org/>

CALL FOR PAPERS

MultiMatch Workshop on 'Innovative Approaches for Searching and Using Cultural Heritage Information for Learning and Teaching'

Berlin, 28 November 2007

The MultiMatch project is organising a workshop on Innovative Approaches for Searching, Collecting, Organizing and Using Multilingual and Multimedia Information for Learning and Teaching, with particular focus on the Cultural Heritage domain. The workshop will be held as a satellite event at Online Educa Berlin 2007. This conference brings together top professionals and experts in technology-supported learning and training of all sectors.

The aim of the MultiMatch sponsored workshop is to converge all stakeholders (educational end users, content providers and technology/service providers) in order to share knowledge and experiences gained in developing, serving and using cultural heritage resources in an e-learning context. The workshop moderator will be Professor Fredrik Truyen, Coordinator ICT for Humanities and Social Sciences at Leuven University and responsible for the MediaLab Institute for Cultural Studies.

Submissions are invited in the form of full and short papers or poster proposals. Topics of interest include, but are not limited to:

- emerging tools/technologies to facilitate cultural heritage education
- using multilingual / multimedia digital resources for learning and teaching
- multimedia applications and virtual reality
- web-based learning
- knowledge discovery in cultural heritage data
- multimedia content analysis and understanding

MultiMatch is developing a search engine offering specialised multilingual search and browse functionality on cultural heritage content from diverse audiovisual sources (including museums, archives and libraries and the web). The first system prototype will be presented during the workshop.

Important Dates

Deadlines for submission of full papers: 31 August 2007

Deadline for submission of short papers and poster proposals: 15 September 2007

Notification to authors: 30 September 2007

Final version of paper: 31 October 2007

The workshop proceedings will be published online.

More information:

http://project.alinari.it/diss-publish/mm_educaberlin.php

CALL FOR PARTICIPATION

GRIDs@Work 2007: Joint European Union - China GRID Days

Beijing, 28 October - 2 November 2007

Further to the success of the third GRID Plugtests in 2006 with some 200 participants from many different countries, ERCIM, INRIA and the European Telecommunications Standards Institute (ETSI) are organizing this year an enhanced GRIDs@work event in Asia. The 2007 GRIDs@work event is composed of a series of workshops and tutorials and the fourth Grid Plugtests.

The Grid Computing Opportunity

Grid computing offers a model for solving massive computational problems using a large number of computers arranged as clusters, embedded in a distributed telecommunications infrastructure. The goals of the event are to bring together Grid researchers, industrials and users; to learn through the Grid users experience about the future features needed for the Grid platform; learn how to best program Grid aware applications; and get important feedback on the deployment and interoperability of Grid applications on various Grid.

Workshops

The event includes the following workshops:

- the second EchoGRID workshop on European-Chinese collaboration in Grid research and technologies
- a GridCOMP workshop and tutorial on designing and implementing a component-based framework suitable to support the development of efficient Grid applications.

Grid Plugtests

This 4th Grid Plugtests will consist of two contests running in parallel: the N-Queens Contest and the FlowShop Contest. During the N-Queens counting challenge, Grid users will try to beat other participants' programs in real-time counting solutions to the N-Queens problem, a constraint satisfaction problem to place N queens on an $N \times N$ 'chess board'. Registration to the Grid Plugtests is free of charge.

Who Should Participate

- researchers and engineers interested in Grid programming techniques
- researchers and engineers who want to tackle the N-Queens world record during the Plugtests event
- software- and hardware vendors, service providers and end-users of various industrial fields.

The event is supported by the Sino-French Laboratory for Computer Science, Automation and Applied Mathematics (LIAMA).

More information:

<http://www.etsi.org/plugtests/Upcoming/GRID2007/GRID2007.htm>

Research Programme on Video and Computer Games at NTNU

The Norwegian University of Science and Technology has established the country's first research programme on video and computer games. The programme complements the country's existing higher education offerings in computer game design and technology, and has as its goal to coordinate and stimulate new research on various aspects of computer games, both within the university itself and among the half-dozen universities and university colleges that offer computer-game related studies. In addition, the programme will work to strengthen the link between the industry and academia with a focus on innovation of computer games. Currently, the programme has established collaborative agreements with Telenor R&I, Arm Norway, FunCom, TellU, Nordisk Film A/S Interactive (the distributor of Sony Playstation in Norway) and Microsoft.

Merlin Forum Launched

The Merlin Global Software Engineering Forum, launched in May 2007, aims at sharing experience and knowledge between companies and learning from each other in the global context. The forum's mission is to create active cooperation between partners from industry and research organisations in the field of global systems and software engineering. It aims to help researchers stay connected with the real life challenges companies face in global product development and on the other hand to help companies stay informed about the latest technologies, tools, and techniques for global development.

The Merlin Global Software Engineering Forum was created as one result of the Merlin (Embedded Systems Engineering in Collaboration, 2004-2007) project. The project addressed the increasing demand to find and discover new efficient ways to support collaborative embedded systems development. Today embedded systems need to be developed globally in collaboration with partners, such as subcontractors, third party developers, and also with in-house developers.

The Merlin Global Software Engineering Forum focuses on facilitating co-operation by arranging regular meetings, featuring the hot topics of systems and software engineering in collaboration. The forum is the vital network that aims to create a base for new and innovative solutions for global systems and software engineering challenges. Active co-operation with industry guarantees also keeping the research strategies up to date and in right direction. Active knowledge gathering and sharing lessons learned as well as dissemination facilitate also launching new projects or creating innovative project ideas, for example. Furthermore, Merlin Forum helps you while searching either industrial or research partners for your needs.

Free online registration to the forum is now open. Also, searching for topics for the next meeting has been started. What is your opinion on today's hottest topic? Do you have suggestions for interesting keynote speakers for future meetings? Or do you have some other ideas on how to develop Merlin forum? For more information, please visit Merlin Forum web-site.

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

Tim Berners-Lee Appointed Member of the Order of Merit by Queen Elizabeth II

Queen Elizabeth II appointed Tim Berners-Lee, W3C Director and inventor of the World Wide Web, to be a member of the Order of Merit. Founded in 1902, the Order of Merit is an honour conferred by the sovereign of the United King-



Tim Berners-Lee.

dom to individuals for "exceptionally meritorious service," usually in the arts, learning, literature and sciences. Twenty four individuals plus foreign recipients may hold the honor at one time. "Awards such as this are for public service, a service which in this case has been largely carried out by the W3C. All those involved in consortium activity should feel recognized by this acknowledgment of the importance of W3C's work," said Berners-Lee.


Second Spanish Conference on Informatics

The Second Spanish Conference on Informatics (CEDI 2007) is to take place between 11-14 September, 2007 in Zaragoza, Spain. The main goal is to join the Spanish informatics research community in order to demonstrate scientific advances, discuss specific problems and improve the visibility of this area in Spain, while emphasizing its role in the age of the information society.

CEDI has a 'federated conference' format, joining several more specific symposia in areas including Artificial Intelligence, Software Engineering and Databases, Programming Languages, Parallelism and Computer Architectures, 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 conference will organize a number of invited talks, round tables, social events and a gala dinner, at which several national awards in Computer Science will be announced: the 'José García Santesteban Award' for the most outstanding professional career, the 'Aritmel Award' for the researcher developing the most significant scientific contributions to informatics engineering, and the 'Mare Nostrum' and 'Ramón Llull' Awards to companies and institutions developing research activities.

More than 1000 researchers are expected to attend the conference. The event will be funded by the Spanish Ministry of Education and Science, and is sponsored by SpaRCIM. The conference is organized by the Department of Informatics and Systems Engineering at the University of Zaragoza. It will be chaired by Víctor Viñals Yúfera from the University of Zaragoza, with Juan José Moreno Navarro (UPM and SpaRCIM director) as Program Chair.

<http://www.congresocedi.es/2007/>



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



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



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



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



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



Centrum voor Wiskunde en Informatica

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



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



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

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



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



Fraunhofer Gruppe
Informations- und Kommunikationstechnik

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



INRIA

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



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



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



Universitas Varsoviensis
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
Tel: +44 1235 44 6103, Fax: +44 1235 44 5147
http://www.scitech.ac.uk/



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



SWEDISH
INSTITUTE OF
COMPUTER
SCIENCE

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



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



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



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

Order Form

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

Send, fax or email this form to:
ERCIM NEWS
2004 route des Lucioles
BP 93
F-06902 Sophia Antipolis Cedex
Fax: +33 4 9238 5011
E-mail: office@ercim.org

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

I wish to subscribe to the

printed edition online edition (email required)

Name:

Organisation/Company:

Address:

Postal Code:

City:

Country:

E-mail: