


Number 74, July 2008

ERCIM



NEWS

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

Special:

Supercomputing at Work

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

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

Editorial Board:

Central editor:

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

Local Editors:

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

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

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

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

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

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

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

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

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

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

Luxembourg: Patrik Hitzelberger (hitzelbe@lippmann.lu)

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

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

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

Sweden: Kersti Hedman (kersti@sics.se)

Switzerland: Harry Rudin (hrudin@smile.ch)

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

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

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

Contributions

Contributions must be submitted to the local editor of your country

Copyright Notice

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

Advertising

For current advertising rates and conditions, see

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

ERCIM News online edition

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

Subscription

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

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

Next issue:

October 2008, Special theme: Safety-Critical Software

Cover illustration:

Massively parallel simulations of aircraft wake instabilities

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

Credits:

Philippe Chatelain, Michael Bergdorf, Diego Rossinelli,

Petros Koumoutsakos, ETH Zurich, Switzerland

Alessandro Curioni, Wanda Andreoni,

IBM Zurich Research Laboratory, Switzerland.

Acknowledgments:

IBM T.J. Watson Research Center

Swiss Supercomputing Center.

Knocking at Petaflop's Door

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

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

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

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

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



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

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

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

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

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

Thomas Lippert

2 Editorial Information

KEYNOTE

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

JOINT ERCIM ACTIONS

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

THE EUROPEAN SCENE

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

SPECIAL THEME

Introduction to the Special Theme

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

31 Microstructural Finite Element Analysis of Human Bone Structures

by Peter Arbenz and Ralph Müller

32 Using Desktop Grids to Securely Store e-Health Data

by Jesus Luna, Manolis Marazakis and Marios D. Dikaiakos

34 Supercomputing in Clinical Practice

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

[Large-Scale Simulation](#)

35 Large-Scale Immune Models and Visualization

by Dimitri Perrin and John Burns

37 High-Performance Computing for Modelling Bacterial Communities

by James T. Murphy, Ray Walshe and Marc Devocelle

38 Harnessing the Power of Supercomputing

by Richard Blake

40 Fluids and Supercomputers: The Billion Particle Era

by Petros Koumoutsakos

41 Supersonic Flow Simulation on Emulated Digital Cellular Neural Networks

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

43 High-Performance Computing of Multiphase Flow

by Bipin Kumar, Yan Delaure and Martin Crane

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

by Dimitrios S. Nikolopoulos

R&D AND TECHNOLOGY TRANSFER

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

by Kashif Din

47 Realizing Ambient Assisted Living Spaces with the PERSONA Platform

by Francesco Furfari and Mohammad-Reza Tazari

49 Medical Record Keeping Made Visually Accessible

by Harry Rudin

50 Games on Networks

by Jacek Miękisz

51 Carrot²: Making Sense of the Haystack

by Stanisław Osiński and Dawid Weiss

52 Quantifying WiMAX Performance

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

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

by Alexander Romanovsky

EVENTS

56 LREC 2008 - The Language Resources and Evaluation Conference

by Nicoletta Calzolari and Khalid Choukri

57 DELOS Summer School on Preservation

by Vittore Casarosa

58 The Fuschl Conversations – A Central Activity of Systems Sciences

by Gerhard Chroust

58 Announcements

62 In Brief

Next issue: October 2008

Next special theme: Safety-Critical Software

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

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

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

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

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

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

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

Link:

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

Please contact:

Kashif Din

Österreichisches Forschungs- und
Prüfzentrum Arsenal GmbH, Austria

Tel: +43 50550 6639

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

Realizing Ambient Assisted Living Spaces with the PERSONA Platform

by Francesco Furfari and Mohammad-Reza Tazari

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

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

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

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

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

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

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

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

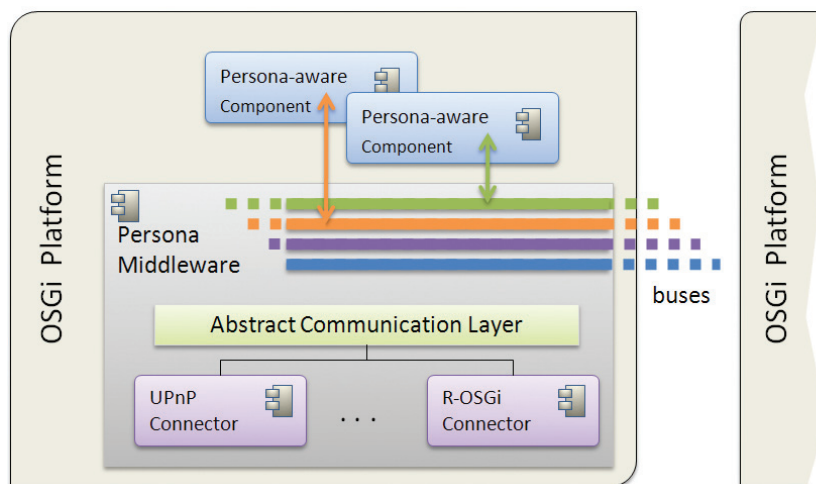


Figure 1: Instances of PERSONA middleware (peers) connected by virtual buses.

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

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

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

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

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

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

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

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

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

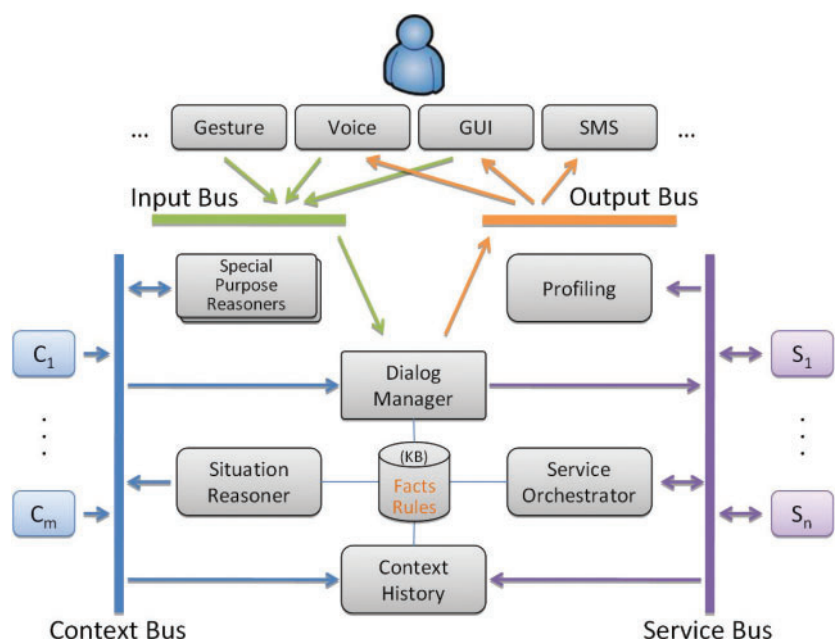


Figure 2: Conceptual architecture of PERSONA.