

# ERCIM



# NEWS

European Research Consortium  
for Informatics and Mathematics  
[www.ercim.eu](http://www.ercim.eu)

Special theme:

# Cloud Computing

Platforms,  
Software  
and Applications

## Also in this issue:

### *Keynote:*

Cloud Computing - The Next  
Big Thing?

*by Keith Jeffery, Burkhard  
Neidecker-Lutz, Lutz Schubert,  
and Maria Tsakali*

### *R&D and Technology Transfer:*

Buiding Discrete Spacetimes  
by Simple Deterministic  
Computations

*by Tommaso Bolognesi*

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

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

#### **Editorial Board:**

*Central editor:*

*Peter Kunz, ERCIM office ([peter.kunz@ercim.eu](mailto:peter.kunz@ercim.eu))*

*Local Editors:*

*Austria: Erwin Schoitsch, ([erwin.schoitsch@ait.ac.at](mailto:erwin.schoitsch@ait.ac.at))*

*Belgium: Benoît Michel ([benoit.michel@uclouvain.be](mailto:benoit.michel@uclouvain.be))*

*Denmark: Jiri Srba ([srba@cs.aau.dk](mailto:srba@cs.aau.dk))*

*Czech Republic: Michal Haindl ([haindl@utia.cas.cz](mailto:haindl@utia.cas.cz))*

*France: Bernard Hidoine ([bernard.hidoine@inria.fr](mailto:bernard.hidoine@inria.fr))*

*Germany: Michael Krapp ([michael.krapp@scai.fraunhofer.de](mailto:michael.krapp@scai.fraunhofer.de))*

*Greece: Eleni Orphanoudakis ([eleni@ics.forth.gr](mailto:eleni@ics.forth.gr))*

*Hungary: Erzsébet Csuhaj-Varjú ([csuhaj@sztaki.hu](mailto:csuhaj@sztaki.hu))*

*Ireland: Ray Walsh ([ray@computing.dcu.ie](mailto:ray@computing.dcu.ie))*

*Italy: Carol Peters ([carol.peters@isti.cnr.it](mailto:carol.peters@isti.cnr.it))*

*Luxembourg: Patrik Hitzelberger ([hitelbe@lippmann.lu](mailto:hitelbe@lippmann.lu))*

*Norway: Truls Gjestland ([truls.gjestland@ime.ntnu.no](mailto:truls.gjestland@ime.ntnu.no))*

*Poland: Hung Son Nguyen ([son@mimuw.edu.pl](mailto:son@mimuw.edu.pl))*

*Portugal: Paulo Ferreira ([paulo.ferreira@inesc-id.pt](mailto:paulo.ferreira@inesc-id.pt))*

*Spain: Christophe Joubert ([joubert@dsic.upv.es](mailto:joubert@dsic.upv.es))*

*Sweden: Kersti Hedman ([kersti@sics.se](mailto:kersti@sics.se))*

*Switzerland: Harry Rudin ([hrudin@smile.ch](mailto:hrudin@smile.ch))*

*The Netherlands: Annette Kik ([Annette.Kik@cwi.nl](mailto:Annette.Kik@cwi.nl))*

*United Kingdom: Martin Prime ([Martin.Prime@stfc.ac.uk](mailto:Martin.Prime@stfc.ac.uk))*

*W3C: Marie-Claire Forgue ([mcf@w3.org](mailto: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.eu/> or contact [peter.kunz@ercim.eu](mailto:peter.kunz@ercim.eu)*

#### **ERCIM News online edition**

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

#### **Subscription**

*Subscribe to ERCIM News by sending email to*

*[en-subscriptions@ercim.eu](mailto:en-subscriptions@ercim.eu) or by filling out the form at the ERCIM*

*News website: <http://ercim-news.ercim.eu/>*

#### **Next issue**

*January 2011, Special theme: "Intelligent and Cognitive Systems"*

#### Keynote

## Cloud Computing: The Next Big Thing?

In 2009 the EC DG Information Society and Media, Software and Services convened an expert group on CLOUD Computing, moderated by Burkhard Neidecker-Lutz of SAP Research and Keith Jeffery of ERCIM, with Lutz Schubert of HLRS as rapporteur and Maria Tsakali as the responsible EC official. The report surveys the current situation of CLOUDs being used both privately within an organisation and as a service external to an organisation. It characterises different kinds of CLOUDs (both existing and future) leading to a list of open research issues that need to be addressed.

A cloud representation has commonly been used in ICT to indicate abstraction or virtualization (eg of a network) and it is this very characteristic that Cloud computing possesses. The system details and management are hidden from the end-user enabling easy outsourcing and utilization of resources. CAPEX (Capital Expenditure) is turned into OPEX (Operational Expenditure) reducing capital expenditure or loan interest. ICT is procured on a 'pay as you go' basis. This not only reduces costs for resource maintenance, but also reduces the risk involved in new product inception, which is one of the major attractions of Clouds. The special capability of Clouds thereby rests on the dynamic and potentially unlimited scalability (both up and down, horizontally and vertically).

The Cloud environment can take one or more of several forms: IaaS (infrastructure as a service), PaaS (platform as a service), AaaS (application(s) as a service) or a totally outsourced ICT capability. Cloud capabilities can implicitly be used to improve the energy efficiency of datacenters, thus supporting the "green ICT" agenda.

Along with the advantages come concerns which need to be addressed for Cloud computing to enjoy significant take-up. The technological problems include: security, trust and privacy; lack of standardisation and therefore supplier lock-in; insufficient virtualization to provide real hiding of systems management (especially in resource sharing/failover) although some PaaS offerings such as Google AppEngine are addressing this issue; data movement and management; programming and system models to provide the required elasticity; systems / services development methods.

There are also non-technological concerns, mainly business / economic / cost models for Cloud computing (including 'green ICT' aspects) that are robust and realistic; legalistic issues concerning data processing, transmission and storage

in another country or multiple countries and/or using an out-sourced service.

However there are major opportunities for Europe in Cloud computing: large companies – especially but not exclusively the telecommunications industry – could provide Cloud services; development by companies (especially SMEs) of products in an open market in Cloud services matching that in goods, services, human capital and knowledge; provision of business model and legalistic expertise (including ‘green ICT’) to accompany the use of Cloud computing.

Are Clouds the next ‘big thing’ in ICT? Clouds are often compared with GRIDs, SOA, Cluster computing and similar technological approaches of the Future Internet. And indeed, CLOUDs typically comprise aspects from all these areas, thus offering improved capabilities for service offering and management. Through its potential globalisation and on demand utilisation, it also offers new business and legalistic models to cost / benefit of ICT.

Clouds within an organisation permit optimisation of ICT in one datacenter (almost certainly replicated – probably externally - for business continuity) increasing resource utilisation and offering server hibernation or switch off varying with demand. This reduces maintenance / systems administration, capital expenditure and energy consumption. The usual business model is that departments in the organisation buy (through internal accounting) services provided in the Cloud so having better cost-management of their ICT; similarly the ICT department is more efficient.

Clouds external to an organisation permit outsourcing of some or all of its IT to another organisation providing the service (and probably providing such a service to several other organisations – multi-tenancy). The customer organisation concentrates on its primary business and treats ICT as a utility service.

Many of the individual research challenges (both technological and non-technological) found in Cloud computing have been addressed through national R&D programmes and the EC framework programme. However, what is needed is to bring these results together in an integrated – and ideally standardised - framework.

The expert group report requests the EC to support R&D in the technological aspects and to set up the required governance framework for Clouds to be effective in Europe. Subsidiary recommendations include the provision of test-beds, joint collaboration groups across academia and industry, standardisation and open source reference implementation (rather like W3C) and the promotion of open source solutions.

If Europe can address and solve the research issues Cloud computing offers a significant opportunity – both for the ICT industry and for commercial activity utilising Clouds.

*Keith Jeffery,  
Burkhard Neidecker-Lutz,  
Lutz Schubert,  
and Maria Tsakali*



*From left: Burkhard Neidecker-Lutz, Keith Jeffery, Maria Tsakali and Lutz Schubert.*

---

**The final report of the Cloud Computing expert group convened by the European Commission DG Information Society and Media, Software and Services is available for download at <http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-report-final.pdf>**

## 2 Editorial Information

### KEYNOTE

- 2 Cloud Computing: The Next Big Thing?**  
by Burkhard Neidecker-Lutz, Keith Jeffery, Maria Tsakali and Lutz Schubert.

### JOINT ERCIM ACTIONS

- 6 Static Analysis versus Model Checking**  
by Flemming Nielson
- 7 An Infrastructure for Clinical Trials for Cancer – ACGT Project Successfully Terminated**  
by Jessica Michel Assoumou and Manolis Tsiknakis
- 8 ERCIM at SAFECOMP 2010 in Vienna**  
by Erwin Schoitsch
- 8 35th International Symposium on Mathematical Foundations of Computer Science**  
by Vaclav Matias
- 9 IWPSE-EVOL 2010 – International Workshop on Principles of Software Evolution**  
by Anthony Cleve and Tom Mens
- 10 ICT Policy Alignment between Europe and India**  
by Nicholas Ferguson, Ashok Kar and Florence Pesce
- 11 Andrea Esuli Winner of the 2010 ERCIM Cor Baayen Award**

## SPECIAL THEME

### Cloud Computing

coordinated by Frédéric Desprez, Ottmar Krämer-Fuhrmann and Ramin Yahyapour

### Introduction to the Special Theme

#### 12 Cloud Computing

by Frédéric Desprez, Ottmar Krämer-Fuhrmann and Ramin Yahyapour

### Invited articles

#### 14 OpenNebula: Leading Innovation in Cloud Computing Management

by Ignacio M. Llorente and Rubén S. Montero

#### 16 SLA@SOI - SLAs Empowering a Dependable Service Economy

by Wolfgang Theilmann and Ramin Yahyapour

#### 18 BEinGRID Presage of the Cloud

by Daniel Field

#### 20 From XtremOS Grids to Contrail Clouds

by Christine Morin, Yvon Jégou and Guillaume Pierre

### Resource management

#### 22 Interoperability between Grids and Clouds

by Attila Marosi, Miklós Kozlovsky and Péter Kacsuk

#### 23 Open Cloud Computing Interface: Open Community Leading Cloud Standards

by Andy Edmonds, Thijs Metsch, Alexander Pappaspyrou and Alexis Richardson

#### 25 Recent Developments in DIET: From Grid to Cloud

by Frédéric Desprez, Luis Rodero-Merino, Eddy Caron and Adrian Muresan

#### 26 Addressing Aggregation of Utility Metering by using Cloud – The Power Grid Case Study

by Orlando Cassano and Stéphane Mouton

#### 27 Optimization and Service Deployment in Private and Public Clouds

by Máté J. Csorba and Poul E. Heegaard

#### 29 Holistic Management for a more Energy-Efficient Cloud Computing

by Eduard Ayguadé and Jordi Torres

#### 30 A Semantic Toolkit for Scheduling in Cloud and Grid Platforms

by András Micsik, Jorge Ejarque, Rosa M. Badia

### Middleware and platforms

#### 32 Making Virtual Research Environments in the Cloud a Reality: the gCube Approach

by Leonardo Candela, Donatella Castelli, Pasquale Pagano

**33 ManuCloud: The Next-Generation Manufacturing as a Service Environment**

by Matthias Meier, Joachim Seidelmann and István Mezgár

**35 RESERVOIR – A European Cloud Computing Project**

by Syed Naqvi and Philippe Massonet

**36 Managing Virtual Resources: Fly through the Sky**

by Jérôme Gallard and Adrien Lèbre

**38 OW2 ProActive Parallel Suite: Building Flexible Enterprise CLOUDs**

by Denis Caromel, Cédric Dalmasso, Christian Delbe, Fabrice Fontenoy and Oleg Smirnov

**40 FoSII - Foundations of Self-Governing ICT Infrastructures**

by Vincent C. Emeakaroha, Michael Maurer, Ivona Brandic and Schahram Dustdar

**41 Large-Scale Cloud Computing Research: Sky Computing on FutureGrid and Grid'5000**

by Pierre Riteau, Maurício Tsugawa, Andréa Matsunaga, José Fortes and Kate Keahey

**43 elasticLM – Software License Management for Distributed Computing Infrastructures**

by Claudio Cacciari, Daniel Mallmann, Csilla Zsigri, Francesco D'Andria, Björn Hagemeier, Angela Rimpl, Wolfgang Ziegler and Josep Martrat

**Applications**

**44 Enabling Reliable MapReduce Applications in Dynamic Cloud Infrastructures**

by Fabrizio Marozzo, Domenico Talia and Paolo Trunfio

**46 Considering Data Locality for Parallel Video Processing**

by Rainer Schmidt and Matthias Rella

**47 Online Gaming in the Cloud**

by Radu Prodan and Vlad Nae

**49 Mastering Data-Intensive Collaboration and Decision Making through a Cloud Infrastructure**

by Nikos Karacapilidis, Stefan Rüping and Isabel Drost

**50 ComCert: Automated Certification of Cloud-based Business Processes**

by Rafael Accorsi and Lutz Lowis

**R&D AND TECHNOLOGY TRANSFER**

**52 Buiding Discrete Spacetimes by Simple Deterministic Computations**

by Tommaso Bolognesi

**54 Improving the Security of Infrastructure Software using Coccinelle**

by Julia Lawall, René Rydhof Hansen, Nicolas Palix and Gilles Muller

**55 Teaching Traffic Lights to Manage Themselves ... and Protect the Environment**

by Dirk Helbing and Stefan Lämmer

**56 A New Approach to the Planning Process makes Huge Savings for the Railway Sector**

by Malin Forsgren and Martin Aronsson

**57 Fast Search in Distributed Video Archives**

by Stephan Veigl, Hartwig Fronthaler and Bernhard Strobl

**58 Meeting Food Quality and Safety Requirements with Active and Intelligent Packaging Techniques**

by Elisabeth Ilie-Zudor, Marcell Szathmári, Zsolt Kemény

**EVENTS**

**60 CLEF 2010: Innovation, Science, Experimentation**

by Nicola Ferro

**61 Announcements**

**IN BRIEF**

**63 Sylvain Lefebvre Winner of the Eurographics Award 2010**

**63 Peter Bosman wins Best Paper Award at Genetic and Evolutionary Computation Conference 2010**

**63 W3C UK and Ireland Office Moves to Nominet**

**63 Mobilize your Apps!**

## Static Analysis versus Model Checking

by Flemming Nielson

*The second annual meeting of the ERCIM Working Group on Models and Logics for Quantitative Analysis (MLQA) took place on Friday July 9th 2010 as part of the Federated Logic Conference (FLoC) organized by the School of Informatics at the University of Edinburgh in Scotland. It was attended by more than 30 researchers, from senior researchers to PhD students, and was one of the best attended satellite events taking place at FLoC.*

The meeting focused on the interplay between static analysis and model checking for verifying and validating IT Systems. Professor Bernhard Steffen, who was perhaps the first researcher to suggest that many static analysis problems can be encoded as model checking problems, gave an overview of how these results emerged two decades ago and the clarity they brought to understanding complex data flow analysis problems and how best to implement them. Professor Steffen's overview was supplemented by Professor Flemming Nielson, who presented recent results showing that state-of-the-art techniques for static analysis, particularly Computation Tree Logic (CTL), can also be used for model checking. This suggests that static analysis and model checking are more similar than classical wisdom would suggest and that algorithmic techniques should be transferred between fields.

Abstraction techniques, lying at the core of static analysis, have for many years been used to reduce the state explosion usually incurred in model checking by introducing uncertainty to cut down on the amount of detail incorporated. In his talk on partial models and software model checking, Dr Arie Gurfinkel gave an overview of a symbolic model checker that allows partial models to interact with the Counterexample-Guided Abstraction Refinement (CEGAR) framework that is an effective technique for gradually increasing the size of models in case a property cannot be established or refuted. This was followed by recent results on Three-Valued Abstraction-Refinement (TVAR), by professor Orna Grumberg, where a third truth value is used to represent the uncertainty, covering examples of its use and a study of the kind of properties that can be verified.

The quantitative dimension focused on validating stochastic properties on systems with probabilistic choices possibly also including nondeterminism. A main challenge is how to extend the abstraction techniques, which work so well in the discrete dimension, to be able to deal with probabilities. Professor Marta Kwiatkowska showed how to use ideas from two-player games to achieve this. Analysing programs mixing probabilities and nondeterminism in a parametric way poses serious challenges, discussed in the presentation by Professor Joost-Pieter Katoen, that showed how to generalize Hoare's axiomatic approach to work with distributions by constructing invariants for linear probabilistic programs.

The challenges of analysing nondeterministic iterative programs were outlined by Dr David Monniaux, whose presentation focused on using techniques other than the abstraction and approximation techniques of Abstract Interpretation, one of the key approaches to static analysis. Instead linear programming techniques and methods from game theory were adapted to construct suitable abstractions. The final presentation, by Dr Michael Huth, "turned the problem upside down" by moving the emphasis from the a posteriori validation of IT Systems to a more holistic approach to the construction of systems guaranteed to live up to quantitative expectations.

In a final business meeting a steering committee was formed for planning the upcoming activities of the working group and for creating links with other communities of researchers sharing in part the vision of MLQA, in particular the community on Quantitative Aspects of Programming Languages (QAPL). The steering committee is comprised of Flemming Nielson (Technical University of Denmark), Diego Latella (ISTI-CNR in Pisa), Joost-Pieter Katoen (RWTH Aachen), Herbert Wiklicky (Imperial College London), Erik de Vink (Eindhoven University of Technology) and Catuscia Palamidessi (INRIA and Ecole Polytechnique).

For more information please consult the MLQA wiki where most of the presentations are available and where details of the next meeting in 2011 will be posted as well as the plans for further stimulating interaction among members of the working group.

**Link:** <http://wiki.ercim.eu/wg/MLQA/>

**Please contact:**

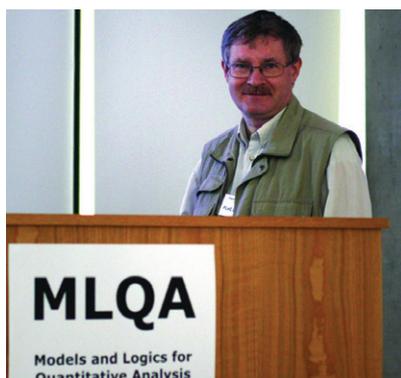
Flemming Nielson

ERCIM MLQA Working Group coordinator

DTU (Technical University of Denmark) Informatics, Denmark

E-mail: [mlqa@imm.dtu.dk](mailto:mlqa@imm.dtu.dk)

*The chairman of MLQA, Professor Flemming Nielson, opens the meeting.*



*MLQA workshop participants.*

# An Infrastructure for Clinical Trials for Cancer – ACGT Project Successfully Terminated

by Jessica Michel Assoumou and Manolis Tsiknakis

*During the last four and one-half years, the EU-funded ACGT project (Advancing Clinico-Genomic Trials on cancer: Open Grid Services for improving Medical Knowledge Discovery) managed by ERCIM, has been developing methods and systems for improved medical knowledge discovery and understanding through the integration of biomedical information.*

The ACGT project vision has been rooted in the realization that information arising from post-genomics research and genetic and clinical trials is rapidly providing the medical and scientific community with new insights, answers and capabilities when combined with advances in high-performance computing and informatics. The objective of the ACGT project has thus been the provision of a unified technological infrastructure which facilitates the seamless and secure access and analysis of multi-level clinico-genomic data enriched with high-performing knowledge discovery operations and services. Biomedical data and information that have been considered include clinical information relating to tissues, organs or personal health-related information, but also information at the level of molecules and cells, as acquired from genomics and proteomics research.

During the course of its life, the project has defined a detailed architectural blueprint and has developed, tested and validated a range of technologies, such as:

- new, domain-specific ontologies, built on established theoretical foundations and taking into account current initiatives, existing standard data representation models, and reference ontologies
- innovative and powerful data exploitation tools, for example multi-scale modelling and simulation, considering and integrating from the molecular to the systems biology level, and from the organ to the living organism level
- standards for exposing the properties of local sources in a federated environment
- a biomedical grid infrastructure offering seamless mediation services for sharing data and data-processing methods and tools
- advanced security tools including anonymisation and pseudonymisation of personal data according to European legal and ethical regulations
- a 'Master Ontology on Cancer' and standard clinical and genomic ontologies and metadata for the semantic integration of heterogeneous databases
- an ontology based 'Trial Builder' for helping to easily set up new clinico-genomic trials, to collect clinical, research and administrative data, and to put researchers in the position to perform cross trial analysis
- data and literature mining services in order to support and improve complex knowledge discovery processes.

## Pilot Trials

The technological infrastructure has been validated in a concrete setting of advanced clinical trials on cancer. The project has targeted two major cancer diseases: breast cancer (BRCA) and paediatric nephroblastoma (PN). The Trastuzumab Optimization Trial in Breast Cancer (TOP), a trial which aims at identifying molecular markers that predict response/resistance to one of the most commonly administered chemotherapies in breast cancer, is one of those pilot trials for ACGT. TOP has been selected as demonstration at the final project review to illustrate several procedures and tools set up by ACGT, for example:

- how the trial data is introduced and analyzed in the ACGT infrastructure to identify the targeted biomarkers
- the importance of the ACGT Master Ontology for the semantic integration of heterogeneous data (clinical, imaging, genomic, proteomic, etc).
- how tools developed within ACGT can facilitate the identification of predictive markers of response/resistance for anthracyclines chemotherapy using microarray-based gene expression profiling as well as genotyping technology
- advances made by the in silico oncology working group. This group has evaluated the reliability of in silico modelling as a tool for assessing alternative cancer treatment strategies; especially in the case of combining and utilizing mixed clinical, imaging and genomic/genetic information and data.

## Outlook

Although the ACGT project is officially ending, the excellent research partnerships developed during the project will continue. The vision of becoming a pan-European voluntary network connecting individuals and institutions while enabling the sharing of data and tools and thus creating a European-wide web of cancer clinical research has been well advanced. The project has developed long lasting partnerships with some of the major stakeholders in the European Cancer Research arena, including ECCO (European Cancer Organisation), BIG (Breast International Group), SIOPE Europe (The European Society for Paediatric Oncology) and the European Clinical Research Infrastructures Network (ECRIN). Building upon the technologies, procedures and knowledge generated by the project, several ACGT partners – jointly with such important end user groups – are about to enter the second phase of implementation. This has been made possible through additional funding from the EU research programmes (both HEALTH and ICT). From legal and ethical aspects to purely ICT solutions, ACGT legacy will be seen under several EU co-financed projects such as the European Network for Cancer Research in Children and Adolescents (ENCCA) and CONTRACT (Consent in a Trial and Care Environment) project that will support the platform created by the ACGT.

In parallel to these developments, the primary instruments for ACGT's collaborative exploitation are well developed. The Center for Data Protection (CDP) has been established and is already actively engaged in service provision. Also the STaRC Initiative has grown into maturity. STaRC is intended to be a 'Study, Trial and Research Centre' that will exploit clinically relevant aspects of ACGT. The concept behind STaRC has received significant recognition and support from patient organisations and patient support groups as well as

from some regional governments. The activities for its official initiation are almost complete.

ACGT was an Integrated Project (IP) funded in the 6th Framework Programme of the European Union under the Action Line "Integrated biomedical information for better health". The project has been carried out by 26 organisations and institutes from academia and industry including the ERCIM members ICS-FORTH (scientific coordinator), INRIA, Computer Architecture Department at the University of Malaga, Technical University of Madrid (both members of SpaRCIM), and Fraunhofer Gesellschaft.

**Link:** <http://eu-acgt.org/>

**Please contact:**

Jessica Michel Assoumou, ERCIM office, France  
E-mail: [jessica.michel@ercim.eu](mailto:jessica.michel@ercim.eu)

Manolis Tsiknakis, FORTH-ICS, Greece  
E-mail: [tsiknaki@ics.forth.gr](mailto:tsiknaki@ics.forth.gr)

## ERCIM at SAFECOMP 2010

by Erwin Schoitsch

ERCIM sponsored the 29th International Conference on Computer Safety, Reliability and Security, SAFECOMP, which took place from 14-17 September 2010 at the Schönbrunn Palace Conference Center in Vienna, Austria. More than 100 experts from 17 countries, including US, Korea, Japan and Brazil, attended and contributed to the conference, workshops and exhibition. The conference was organized by the Austrian Computer Society (OCG), the Austrian Institute of Technology (AIT) and EWICS TC7 (European Workshop on Industrial Computer Systems, TC7, Reliability, Safety and Security). The conference and programme chair was Erwin Schoitsch from AARIT/AIT.

SAFECOMP is a specialized and leading international conference on Computer Safety, Reliability and Security, established 1979 by EWICS TC7 (Purdue Europe). The 29th Conference had an attractive programme dealing with various aspects of critical embedded systems engineering, system analysis, testing, system modelling, design, development, verification and validation, standards, safety cases and certification, and application related aspects of safety, reliability and security in automotive, aerospace, railways, critical infrastructures (smart grids).

The first day was dedicated to workshops where the ERCIM Working Group on Dependable Embedded Systems co-organized the ERCIM/ DECOS/MOGENTES Workshop on Dependable Embedded Systems, chaired by Amund Skavhaug and Erwin Schoitsch. The papers and presentations will be published by ERCIM. ERCIM was presented in the workshop by Erwin Schoitsch and with booth in the conference exhibition.

Three invited keynotes, one at each day, were the highlights of the conference:



Joint ERCIM-OCG booth at SAFECOMP (left) and the conference venue Schönbrunn Palace in Vienna.

- "System of Systems Challenges" by Hermann Kopetz - Vienna University of Technology, Austria
- "Murphy Was An Optimist", by Kevin Driscoll - Honeywell Laboratories, USA
- "Process Control Security: go Dutch! (united, shared, lean and mean)", by Eric Luijff - TNO, The Hague, The Netherlands.

The proceedings will be published by Springer in the Lecture Notes in Computer Science series 6351.

A "best paper award" was granted to Bernd Fischer, and his co-authors Nurlida Basir (both University of Southampton, UK) and Ewen Denney, NASA Ames Research Center, USA for their paper "Deriving Safety Cases for Hierarchical Structure in Model-Based Development".

The next conference - SAFECOMP 2011 - will be held in Naples, Italy from 19-21 September 2011.

**Link:** <http://www.ocg.at/safecomp2010/>

## 35th International Symposium on Mathematical Foundations of Computer Science

by Vaclav Matias

In the 2010, the 35th International Symposium on Mathematical Foundations of Computer Science (MFCS 2010), sponsored by ERCIM and 19th EACSL Annual Conferences on Computer Science Logic (CSL 2010) were federated and organized in parallel. The scientific program of MFCS & CSL 2010 was further enriched by twelve satellite workshops on more specialized topics. The symposium was attended by more than 350 participants from 38 countries and five continents.

The main conferences and their satellite events were hosted by the Faculty of Informatics, Masaryk University, Brno, Czech Republic, on 21-29 August 2010. Masaryk University



ERCIM booth at MFCS.

if a founding member of CRCIM, the Czech member institution of ERCIM.

The series of MFCS symposia, organized in rotation by Poland, Slovakia, and the Czech Republic since 1972, encourage high-quality research in all branches of theoretical computer science. Their broad scope provides an opportunity to bring together researchers who do not usually meet at specialized conferences.

Computer Science Logic (CSL) is the annual conference of the European Association for Computer Science Logic (EACSL). The conference is intended for computer scientists whose research activities involve logic, as well as for logicians working on issues significant for computer science.

**Link:** <http://mfcs12010.fi.muni.cz/>

## IWPSE-EVOL 2010 – International Workshop on Principles of Software Evolution

by Anthony Cleve and Tom Mens

*The sixth annual workshop of the ERCIM Working Group of Software Evolution took place in Antwerp, Belgium, 20-21 September 2010, under the auspices of the ERCIM Working Group on Software Evolution. The event gathered theorists and practitioners to present and discuss the state-of-the-art in research and practice on automated software evolution.*

For the second year in a row, the ERCIM Working Group on Software Evolution jointly organized its annual workshop on software evolution (EVOL) together with the International Workshop on Principles of Software Evolution (IWPSE). This year's edition focused on Automated Software Evolution, in order to remain in sync with the international conference with which IWPSE-EVOL 2010 was co-located, namely ASE 2010, the IEEE/ACM international conference on Automated Software Engineering that celebrated its 25th anniversary edition. The workshop was co-organized by Anthony Cleve (ERCIM Postdoctoral Fellow at INRIA Lille, France), Naouel Moha (INRIA Rennes, France) and Andrea Capiluppi (University of East London, UK).

IWPSE-EVOL was the most successful workshop of ASE 2010 with 28 participants originating from 12 different countries across all continents. The workshop attracted 31 submissions. To maintain the high quality standards of our workshop, only 13 of these submissions were accepted for presentation and publication, after a rigorous peer review by at least three different members from the programme committee. All accepted papers have been published in the ACM International Conference Proceedings Series (AICPS). The full list of papers and authors can be seen on the workshop's website. In addition, a special issue of Elsevier's Journal of Systems and Software (JSS) dedicated to automated soft-

ware evolution will be associated to IWPSE-EVOL 2010. Paper submission deadline is 30 November 2010.

A best paper award (see photo) was presented to Lile Hattori, Mircea Lungu and Michele Lanza, for their paper "Replaying past changes in multi-developer projects". This paper presents Replay, an interactive tool allowing software developers to replay past software changes at a fine-grained level.



*Anthony Cleve (left), ERCIM fellow and workshop co-organiser, and Lile Hattori, first author of the award-winning paper.*

The workshop also welcomed two invited keynote talks. The first one, by Andrian Marcus (Wayne State University, USA) was entitled "Software is Data Too: How should we deal with it?" In this talk, Andrian Marcus discussed the numerous challenges related to the analysis and management of software data. He argued for new empirical research approaches and for stronger collaboration among researchers and practitioners. The second invited talk, by Massimiliano Di Penta (University of Sannio, Italy), was entitled "Empirical Studies on Software Evolution: Should we (try to) claim Causation?". Massimiliano Di Penta presented his vision of empirical studies on software evolution. He elaborated on the need to integrate new sources of information, to combine several analysis techniques and, last but not least, to carefully interpret the results obtained.

### **Links:**

**Workshop Proceedings:**

<http://portal.acm.org/citation.cfm?id=1862372>

**Workshop website:** <http://soft.vub.ac.be/iwpse-evol/>

**Working Group website:**

<http://wiki.ercim.eu/wg/SoftwareEvolution>

### **Please contact:**

Tom Mens,

ERCIM Software Evolution Working Group chair  
Université de Mons, Belgium

Tel: +32 65 37 3453

E-mail: [tom.mens@umons.ac.be](mailto:tom.mens@umons.ac.be)

## ICT Policy Alignment between Europe and India

by Nicholas Ferguson, Ashok Kar and Florence Pesce

***The ERCIM-led Euro-India SPIRIT project launches the first Information and Communication Technologies (ICT) Experts meeting to discuss ICT policy alignment between Europe and India.***

Renowned ICT experts from the European Union (EU) and from India met in Hyderabad, India in August to deliberate on a joint approach to conduct research and foster innovation across emerging technologies and critical societal applications enabled by ICTs. This endeavour will pave the way for aligning their existing and future policies driving ICT research and recommend a strategy for increasing and supporting joint programmes and projects to address key technological and societal priorities in India and in the European Union.

Organised in three thematic groups – ICTs addressing Societal Challenges, AudioVisual Media & Internet and Emerging Technologies & eInfrastructures – the experts cover a large spectrum of ICT actors - from cutting edge research institutions to dedicated NGOs taking the benefits of ICTs to the masses. Professor Ashok Jhunjhunwala (IIT Madras), Professor T.V. Prabhakar (IIT Kanpur), Professor Krithi Ramamritham (IIT Bombay) and Professor P.J. Narayanan (IIIT Hyderabad) bring academic and research excellence, while Dr. Vinay Deshpande (Encore Software), Dr. Hiranmay Ghosh (TCS Innovation Labs), Subu Goparaju (SET Lab, Infosys Technologies) come on board with pioneering corporate research in ICT and Ms. P.N. Vasanti (Centre for Media Studies) and Ms. Anita Gurusurthy (IT for Change) anchor the deliberations in societal priorities. The European Union experts are a rich and diverse mix – Jim Clarke (Waterford Institute of Technology, Ireland), Pierre-Yves Danet (Orange Lab – France Telecom), Professor Soumitra Dutta (INSEAD, France), Dr. Fabrizio Gagliardi (Microsoft Research, Europe), Professor Thorsten Herfet (University of Saarland, Germany), Dr. Mounib Mekhilef (Ability Europe, UK), Professor Mogens Kuehn Pedersen (Copenhagen Business School, Denmark), Professor Neeraj Suri (Technological University of Darmstadt, Germany) and Ms. Nicole Turbe-Suetens (Distance Expert, France).

After a day of intense discussion, the experts presented key findings at a workshop on Collaborative ICT Research & Innovation between India and the European Union organised by the Euro-India SPIRIT project at eIndia2010, India's largest ICT event. Over 50 participants were offered a comprehensive view of the ICT research and innovation priorities and programmes in the European Union while benchmarking the Indian ICT research objectives and initiatives and offering a vision of the way ahead in this collaborative endeavour.

Speaking at eIndia2010, Vinay Deshpande Encore Software and member of the Emerging Technologies & eInfrastructures WG commented, "Indian ICT can benefit from its interaction with the EU. In terms of policy we can learn from what has been done and the experiences of poli-

cies that have been adopted in the EU relating to e-Governance, social systems and how they can be adapted and even re-aligned for the needs of India. Likewise, the successes of Indian ICT can be transported to the EU so collaboration is a 2-way street."

For its part, Europe is ready to learn from the Indian ICT industry with its world-renowned expertise in software development, especially in e-Governance, social systems and the rural sector. Fabrizio Gagliardi, Microsoft Research and member of the Emerging Technologies & eInfrastructures WG, remarked, "We have identified cloud computing as one of the potential new technologies and trends that can enable better science and improved industrial applications in India, especially with the highly distributed geographical nature of India." Such an approach perfectly illustrates how closer co-operation on emerging technologies like Cloud can bring benefits across the board.

The benefits of EU-Indian ICT collaboration were also underlined with regard to the aging work force and the development of collaborative working environments which can offer a better work-life balance to aged workers while capitalising on their knowledge and reducing stressful factors such as travelling for work. According to Nicole Turbé-Suetens, Distance Expert and member of the Societal Challenges WG, "EU-Indian collaboration is helping people understand how these technologies can enhance their processes and deliver better quality of life, services and help their development"

Similarly, in the Telecommunications field, Pierre Yves Danet, France Telecom, and AudioVisual and Internet WG member believes that "technology must be aligned to ensure interoperability between different services. When setting up a service, particularly a mobile service, it is important that people from India, Europe and around the world can use the same phone and the same services."

Joint EU and Indian ICT initiatives are a catalyst in addressing many of the challenges faced by India today such as improving public healthcare infrastructures and services, bringing efficient governance and seamless transactions to citizens, safeguarding the environment, improving education, learning resources and tools for the masses. Collaborative international research is an essential ingredient to achieve India's ambitions in ICT, at home and abroad. The deliberations in Hyderabad offered fertile ground to strengthen this co-operation and ensure adequate policy measures are put in place.

The Working Groups will continue their discussions on 9 November 2010 and will prepare recommendations to be presented at the Euro-India SPIRIT Workshop to be held in Delhi on 10 November. The workshop will be co-located with the science & technology stakeholders conference organised on 11 and 12 November by the Delegation of the European Union to India.

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

**Please contact:**

Florence Pesce, ERCIM office, France

E-mail: [florence.pesce@ercim.eu](mailto:florence.pesce@ercim.eu)

## Andrea Esuli Winner of the 2010 ERCIM Cor Baayen Award

*Andrea Esuli from ISTI-CNR has been chosen by ERCIM as the winner of the 2010 Cor Baayen Award for a promising young researcher in computer science and applied mathematics.*

During the course of his PhD work and his postdoctoral research, Andrea Esuli has obtained outstanding results by combining cutting-edge speculative reasoning with top-quality technical solutions, and has produced results of unquestionable societal and commercial value.

One of Esuli's distinctive characteristics as a researcher is his broad range of capabilities. He is neither a purely theoretical thinker, nor a mere experimenter, but a mature researcher with the ability to conceive innovative solutions and immediately test them. This is so thanks to a strong theoretical background and to excellent technological abilities. Esuli's activities and results range from academic research to the development of industrial-strength innovative applications. Concerning the former aspect, he must be credited (among others) with being one of the early researchers in the field of sentiment analysis and opinion mining. This field, that is so popular nowadays, was still a tiny niche in 2005 when Andrea published his first paper on the subject, at a highly selective ACM conference. Nowadays, Andrea is recognized as a top player in the field; he was invited to speak on this topic as a panelist at the GWC 2008 conference, even before getting his PhD, which is a rare feat.

One of the results of his PhD research on lexical resources for opinion mining is SentiWordNet, a lexical resource licensed to more than 400 research labs / companies worldwide, and widely considered the reference lexical resource for opinion mining applications. The systems developed in the course of his research on supervised machine learning algorithms, training data cleaning, and active learning for automatic text classification, are now part of the Verbatim Coding System (VCS), a highly successful software system for the analysis of textual answers returned by respondents to questions issued in the context of opinion surveys. VCS, the recipient of one international and one national award, has been licensed to corporate end users whose Customer Relationship Management departments use it in their daily operations, and is now an integral part of the Ascribe(TM) platform marketed by Language Logic (<http://www.languagelogic.info/>), the world leader in the provision of survey management software services. A more recent thread of Esuli's research is on highly efficient similarity search, for which he has developed a novel algorithm based on prefix-permutation indexing. In addition to generating scientific publications, Andrea has turned this algorithm into a working search engine for images (<http://mipai.esuli.it/>) that currently allows image similarity search on CoPhIR, the largest image dataset available for research purposes. Esuli's algorithm allows similarity searches to be conducted on CoPhIR in sub-second response times, a feat currently nei-



*Andrea Esuli.*

ther attained nor approached by competing systems. This algorithm has resulted in a patent submission currently under review by the US Patent and Trademark Office.

Andrea Esuli holds an MSc in Computer Science (2001), an MSc in Computer Science Technologies (2003), and a PhD in Information Engineering (2008), all from the University of Pisa, and all with full marks and "cum laude". From 2002 to 2004 he was a research associate at the Department of Computer Science of the University of Pisa, where he worked in the area of high-performance information retrieval, with particular emphasis on algorithms and data structures for large-scale collaborative text indexing and "query search" processes. Since 2005, he is with the Istituto di Scienza e Tecnologie dell'Informazione, Consiglio Nazionale delle Ricerche, Pisa, where he has carried out research in the areas of text classification, information extraction, and opinion mining. He has a great research record for this stage of his career and has authored some truly influential papers.

### About the Cor Baayen Award

The Cor Baayen Award is awarded each year to a promising young researcher in computer science and applied mathematics. The award was created in 1995 to honour Cor Baayen, the first president of ERCIM and the ERCIM 'president d'honneur'. The award consists of a cheque for 5000 Euro together with an award certificate.

#### 2010 Finalists

The ERCIM Executive Committee has accepted twelve finalists (in alphabetical order):

- Maxime Descoteaux
- Andrea Esuli
- Jose M. Juarez
- Arti Klami
- Pushmeet Kohli
- Claudio Lucchese
- Michael Mavroforakis
- Soren Sonnenburg
- Sven Schewe
- Paschalis Tsiaflakis
- Jeroen Wackers
- Olaf Zimmermann

The winner, Andrea Esuli, was selected by the ERCIM Executive Committee on advice from the ERCIM Advisory Committee.

#### More information:

<http://www.ercim.eu/activity/cor-baayen-award>

# Cloud Com

## Introduction to the Special Theme

by Frédéric Desprez, Ottmar Krämer-Fuhrmann and Ramin Yahyapour

The fast evolution of hardware capabilities in conjunction with fast wide area communication and availability of virtualization solutions is allowing new operational models for information technology. The advent of Cloud computing has resulted in access to a wide range of services. Infrastructure-as-a-Service (IaaS) allows access to large-scale resources like computational power or storage. These large scale platforms, based on huge datacentres, are available to industrial users as well as to scientific communities. IaaS allows the use of a large number of bare machines on which any software stack can be installed. The Platform-as-a-Service (PaaS) model provides the programmer with sets of software elements that can be combined in a scalable way to build large scale applications. Finally, Software-as-a-Service (SaaS) simplifies access to large applications in a remote and seamless way.

Access to these platforms is driven through the use of different services providing mandatory features such as security, resource discovery, virtualization, load-balancing, etc. Platform as a Service as well as Software as a Service thus have an important role to play in the future development of large scale applications. The overall idea is now to consider the whole system, ranging from the resources to the application, as a set of services. Hence, a user application is an ordered set of instructions requiring and making use of services, for example a service of execution.

The variety of service-based platforms and the way in which they are accessed as Cloud services have an important impact

on how applications are designed (ie, the programming model used) as well as how applications are executed (ie, the runtime/middleware system used). While core technologies for this kind of abstraction (virtualization, Grids etc), have existed for many years we are only now beginning to see a broad proliferation and a disruptive technology change in our perception on network-based services.

This special issue of ERCIM news gathers papers from the leading European research activities in this area. Invited contributions highlight the innovations from four key European projects: OpenNebula, SLA@SOI, BEinGRID and XtremOS.

The following papers cover important topics relating to the main challenges in future Cloud platforms: programming and access, resource management, middleware and platforms, and applications.

### Programming and access

Programming and accessing Cloud platforms in a seamless and efficient way is one of the most important challenges for the next decade. Today's applications ported from Grids are still being programmed either as regular batch processing jobs or message passing applications. Neither of these models will be able to scale and perform on these large scale and dynamic platforms. The near future will provide new computational platforms with a large number of processing cores. There is an inherent need to understand how such systems can be efficiently put to use. Existing parallelization models might not be sufficient to cope with this challenge. As such, novel approaches are necessary.

# Computing

## Resource management

Resource management has been the first issue raised over IaaS platforms. Scaling with thousands of virtual machines accessed by thousands of users remains a crucial challenge for these platforms. Existing approaches from Grids and HPC focus largely on application performance. However, the management of future platforms is more likely to become a multi-criteria problem also taking into account energy consumption, cost and quality of service. The allocation and management problem can also be seen from different viewpoints: first, the service provider needs to manage the underlying infrastructure for provisioning; second, the service consumer needs to monitor the portfolio of services she is relying upon. Any efficient resource management system also needs suitable support for monitoring and measurement. The use of service-level agreements has become a common requirement to support business scenarios and provide a clear understanding on the details of service contracts.

## Middleware and platforms

Middleware systems remain the cornerstone of platforms like Clouds. Being able to manage virtual machines efficiently over several Cloud platforms is an important research and development issue. Migration has to be performed efficiently in order to manage dynamic workload and cope with the dynamicity of nodes and applications. Now that PaaS and SaaS platforms are becoming widely available, practical problems like license management issues need to be resolved in order to provide service consumers with the required software environments.

Current research is also addressing questions of interoperability and federation of Cloud platforms. As there are many large Grid infrastructures in operations, their coexistence and collaboration with Clouds is of major importance.

## Applications

Finally, the scope of applications ported over large Cloud platforms will be larger than ever including new applications such as online games and video processing. These new and highly distributed applications will raise new research problems such as quality of service and latency management.

## Please contact:

Frédéric Desprez  
INRIA, France  
E-mail: frederic.desprez@inria.fr

Ottmar Krämer-Fuhrmann  
Fraunhofer Institute for Algorithms and Scientific Computing  
SCAI, Germany  
E-mail: ottmar.kraemer-fuhrmann@scai.fraunhofer.de

Ramin Yahyapour  
Technical University Dortmund, Germany  
E-mail: ramin.yahyapour@tu-dortmund.de

# OpenNebula: Leading Innovation in Cloud Computing Management

by Ignacio M. Llorente and Rubén S. Montero

*OpenNebula is the result of many years of research and development in efficient and scalable management of virtual machines on large-scale distributed infrastructures. Its innovative features have been developed to address the requirements of business use cases from leading companies in the context of flagship European projects in cloud computing. OpenNebula is being used as an open platform for innovation in several international projects to research the challenges that arise in cloud management, and also as production-ready tool in both academia and industry to manage clouds.*

As virtualization technologies mature at an incredibly rapid pace, there is a growing interest in applying them to the data-centre. After the success of cloud computing, companies are seeking reliable and efficient technologies to transform their rigid infrastructure into a flexible and agile provisioning platform. These so-called private clouds allow you to provide IT services with an elastic capacity, obtained from your local resources in the form of Virtual Machines (VM). Local resources can be further combined with public clouds in a hybrid cloud computing setup, thus enabling highly scalable hosting environments.

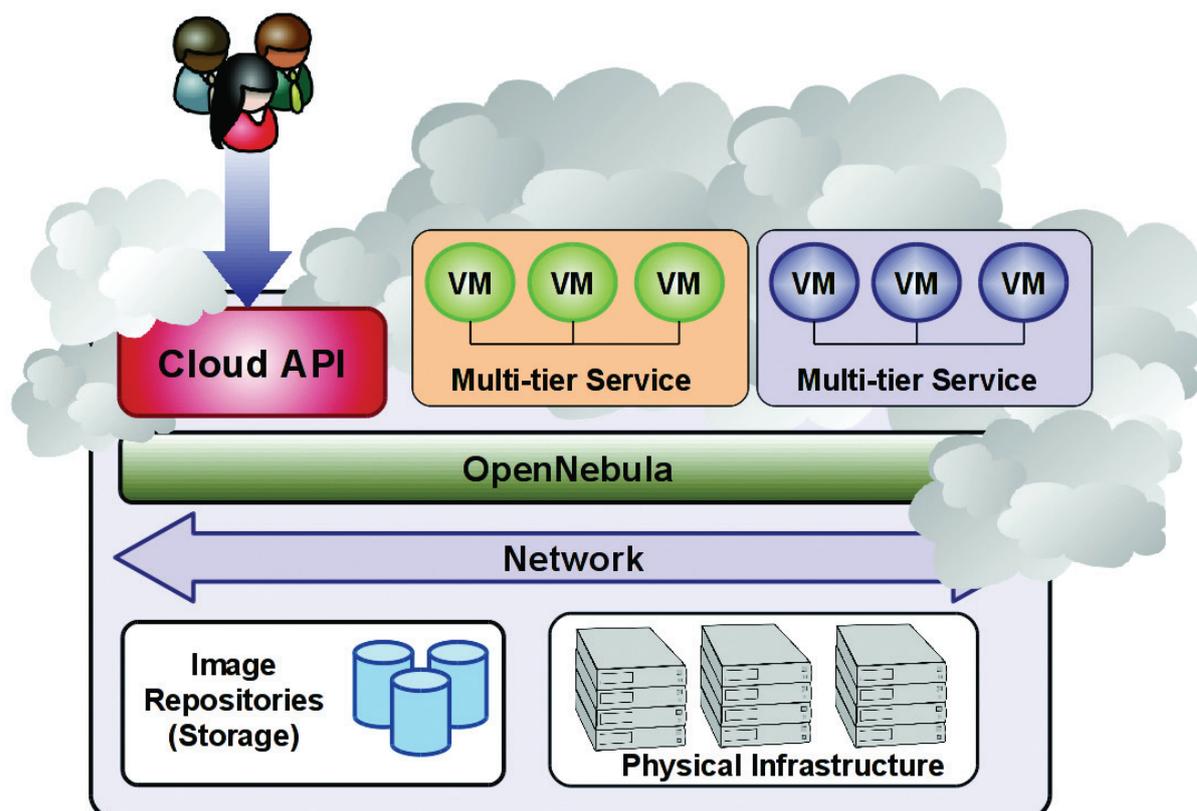
The main component involved in implementing this provision scheme is the Cloud Management Tool, which is responsible for the secure, efficient and scalable management of the cloud resources. A Cloud Management Tool provides IT staff with a uniform management layer across distributed hypervisors and cloud providers; giving infrastructure users the impression of interacting with a single infinite capacity and elastic cloud.

Because no two data centres are the same, building clouds is about integration and orchestration of the underlying infrastructure systems, services and

processes. The Cloud Management Tool should seamlessly integrate any existing security, virtualization, storage, and network solutions deployed in the data-centre. Moreover, the right design and configuration in the Cloud architecture depend not only on the underlying infrastructure but also on the execution requirements of the service workload. The capacity requirements of the virtual machines as well as their level of coupling determine the best hardware configuration for the networking, computing and storage subsystems.

OpenNebula is an open-source Cloud Management Tool that embraces this

*OpenNebula architecture.*



vision. Its open, architecture, interfaces and components provide the flexibility and extensibility that many enterprise IT shops need for internal cloud adoption. These features also facilitate its integration with any product and service in the cloud and virtualization ecosystem, and management tool in the data centre. OpenNebula provides an abstraction layer independent from underlying services for security, virtualization, networking and storage, avoiding vendor lock-in and enabling interoperability. OpenNebula is not only built on standards, but has also provided reference implementation of open community specifications, such as the OGF Open Cloud Computing Interface. This open and flexible approach for cloud management ensures widest possible market and user acceptability, and simplifies adaptation to different environments.

#### From Research Project to Open Software, Community and Ecosystem

OpenNebula was first established as a research project back in 2005 by the Distributed Systems Architecture Research Group at the Complutense University of Madrid. Since its first public release of software in March 2008, it has evolved through open-source releases to a strong user community and now operates as an open source project. The OpenNebula project has a strong commitment with open-source, being one of the few cloud management tools that are available under Apache license. The Apache license allows any cloud and virtualization player to innovate using the technology without the obligation to contribute those innovations back to the open source community.

The OpenNebula technology has matured thanks to an active and engaged community of users and developers. OpenNebula is downloaded several thousands times per month from its site, and the code can be also downloaded from the software repository and from several commercial and open-source distributions. The development is driven by its community in order to support the most demanded features, and by the international research projects funding OpenNebula in order to address the demanding requirements of several business and scientific use cases for cloud computing. OpenNebula has proved to be a production-ready solu-

tion that includes enterprise features such as security, robustness, scalability and performance that many IT shops need for internal cloud adoption, either in scientific or business environments.

Besides an exponential growth in its number of users, there are many projects, research groups and companies building new virtualization and cloud components to complement and to enhance its functionality. These components build the quickly evolving OpenNebula ecosystem, in which related tools, extensions and plug-ins are available from and for the community. Additionally, OpenNebula leverages the ecosystems being built around other popular cloud interfaces, such as Amazon AWS and VMware vCloud.

#### Outlook

OpenNebula is being funded by several Spanish projects, such as NUBA (strategic research program), BIOGRIDNET, Grid4Utility, HPCcloud and MEDIANET. Some of its main enhancements have been developed in RESERVOIR, flagship of European projects in cloud computing research. It is also used as core technology in many new European projects, such as StratusLab, aimed at bringing cloud and virtualization to grid computing infrastructures; BonFIRE, which targets the services research community on Future Internet, with the aim of designing, building and operating a multi-site cloud-based facility to support research across applications, services and systems; and 4CaaS, aimed at creating an advanced PaaS Cloud platform which supports the optimized and elastic hosting of Internet-scale multi-tier applications.

Existing research funding ensures the engineering resources to support and develop OpenNebula and thus to maintain OpenNebula's position as the leading and most advanced open-source technology to build cloud infrastructures. Additionally, C12G Labs is a new start-up that has been created to provide the professional integration, certification and technical support that many enterprise IT shops require for internal adoption. This also contributes to OpenNebula's long term sustainability by ensuring that it is not tied exclusively to public financing.

#### Links:

DSA-Research: [www.dsa-research.org](http://www.dsa-research.org)  
OCCI-WG: [www.occi-wg.org](http://www.occi-wg.org)  
RESERVOIR: [www.reservoir-fp7.eu](http://www.reservoir-fp7.eu)  
OpenNebula: [www.opennebula.org](http://www.opennebula.org)  
StratusLab: [www.stratuslab.eu](http://www.stratuslab.eu)  
BonFIRE: [www.bonfire-project.com](http://www.bonfire-project.com)  
4CaaS: [4caast.morfeo-project.org](http://4caast.morfeo-project.org)  
C12G: [www.c12g.com](http://www.c12g.com)

#### Please contact:

Ignacio M. Llorente,  
Tel: +34 91 3947616  
E-mail: [llorente@dacya.ucm.es](mailto:llorente@dacya.ucm.es)

Rubén S. Montero  
Tel: +34 91 3947538  
E-mail: [rubensm@dacya.ucm.es](mailto:rubensm@dacya.ucm.es)

Distributed Systems Architecture  
Research Group, Complutense  
University of Madrid, Spain

# SLA@SOI - SLAs Empowering a Dependable Service Economy

by Wolfgang Theilmann and Ramin Yahyapour

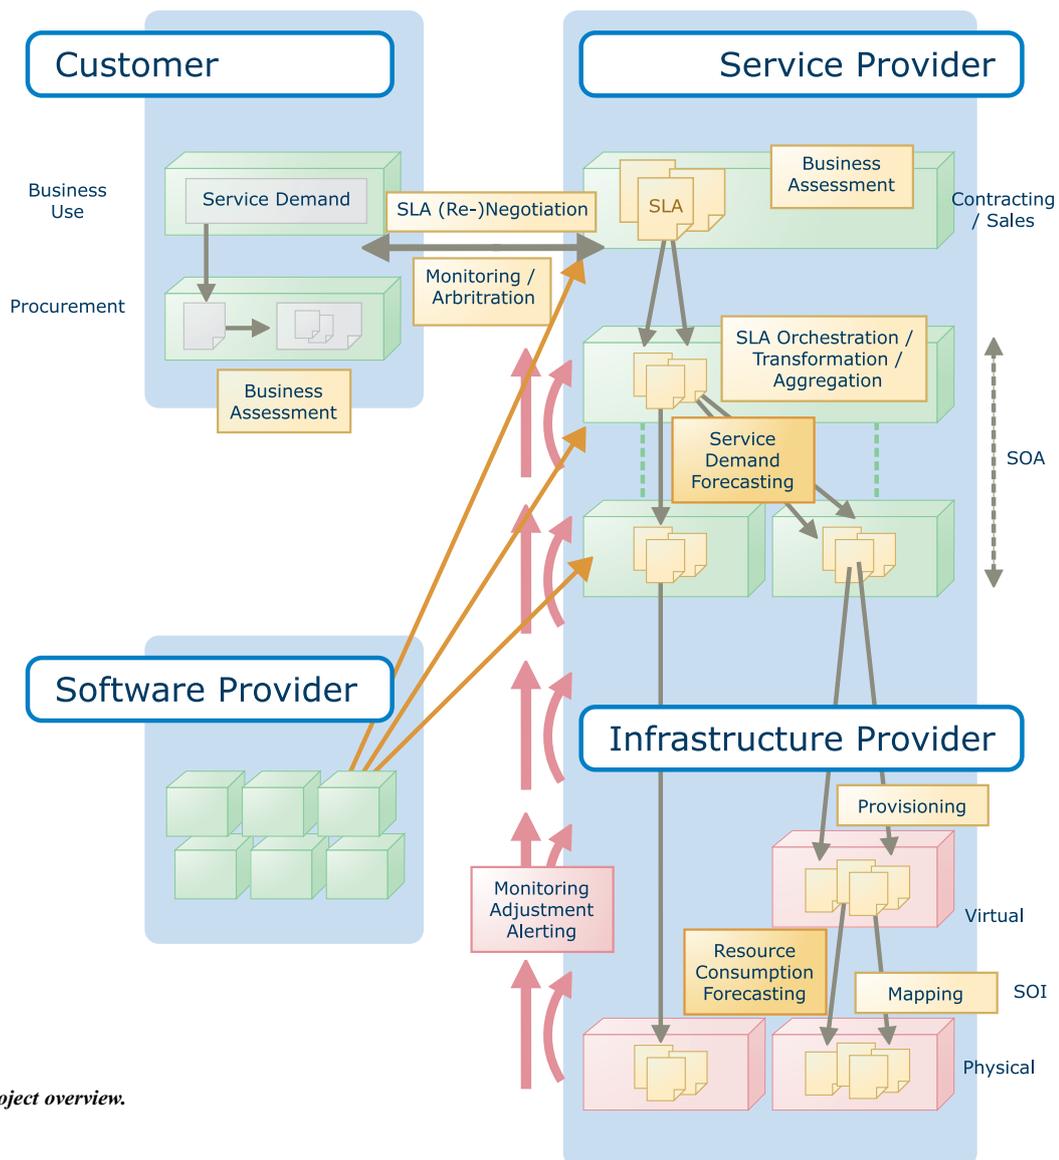
*IT-supported service provisioning has become of major relevance in all industries and domains. The research project SLA@SOI provides a major milestone for the further evolution towards a service-oriented economy, where IT-based services can be flexibly traded as economic goods, ie under well-defined and dependable conditions and with clearly associated costs. Eventually, this will allow for dynamic value networks that can be flexibly instantiated thus driving innovation and competitiveness. SLA@SOI created a holistic view for the management of service level agreements (SLAs) and provides an SLA management framework that can be easily integrated into a service-oriented infrastructure.*

Europe has set high goals in becoming the most active and productive service economy in the world. This is especially true for IT supported services, which have evolved into a common utility which is offered and consumed by many stakeholders. Cloud Computing, for instance, has gained significant attention

and commercial uptake in many business scenarios. With more companies incorporating cloud-based IT services as part of their own value chain, reliability and dependability become crucial factors in managing business. Service-level agreements are the common means to provide the necessary trans-

parency between service consumers and providers.

SLA@SOI is a major European project that addresses the issues surrounding the implementation of automated SLA management solutions on Service Oriented Infrastructures (SOI) and eval-



SLA@SOI project overview.

uates their effectiveness. SLA's are particularly relevant to cloud computing, an increasingly important and relevant deployment model for infrastructure, services or platforms. SLA@SOI allows such services to be described by service providers through formal template SLA's. Once these template SLAs are machine readable, service composition can be established using automatic negotiation of SLAs. Moreover, the management of the service landscape can focus on the existence and state of all necessary SLAs.

A major innovation of SLA@SOI is the multi-layered aspect of the service stack. Typically, a service is dependent on many other services. For example, the offering of a software service requires infrastructure resources, software licenses or other software services. SLA@SOI's SLA framework allows the configuration of complex service hierarchies with arbitrary layers.

#### Architecture

The technical foundation to the functional and business innovations is a highly configurable plug-in-based architecture, supporting flexible deployment options which are incorporated into existing service landscapes. The primarily open source implementation embraces the latest open standards. A harmonized virtualization infrastructure supports private, public and hybrid clouds, whether they use commercial or open source provisioning systems. The accompanying figure illustrates the anticipated SLA management activities throughout the Business/IT stack.

The framework's architecture mainly focuses on separation of concerns, related to SLAs and services on the one hand, and to the specific domain (eg, business, software, and infrastructure) on the other. Service Managers are responsible for all management activities directly related to services. This includes the management of information about available services, supported types of services, as well as their offered functionality and their dependencies. SLA Managers are responsible for all actions that are related to service-level agreements. They are involved in negotiation with customers and are responsible for the planning and optimization of new services that are to be provisioned. Furthermore, they monitor the terms upon which a provider and

customer have agreed and react in case of violations. SLA Managers can negotiate with each other in order to find the best offer for a customer. The provisioning of a service is a joined effort of all SLA Managers and Service Managers involved. In order to support multiple domains with our framework, multiple SLA managers and multiple Service Managers can collaborate inside the framework as well as across framework boundaries. Thereby, each SLA Manager and Service Manager is responsible for SLAs and services of a particular domain. The root of the management hierarchy is the Business Manager component. It is responsible for asserting overall business constraints on the system in order to meet business objectives and for maintaining customer and provider relations..

#### Innovations

SLA@SOI advances the state-of-the-art by realizing an open, powerful and flexible framework for SLA-enabling any service with an arbitrarily complex SLA. The innovations span not just functional concerns but also the business domain, and are constructed on a robust technical foundation. Technically, SLA@SOI will release a comprehensive open-source with reference models and plug-in implementations for common deployment scenarios. From a scientific point of view, it will publish new and evolved algorithms and models as well as contribute to open standards such as Open Grid Forum's WS-Agreement and Open Cloud Computing Interface (OCCI), or DMTF's OVF. The key innovations from a business perspective include a business management suite for automated e-contracting and post-sales management. Domain-specific adoption guidelines including reference templates and models will also be provided, simplifying integration of the framework into existing service landscapes.

#### Use Cases

SLA@SOI is a comprehensive project with a broad scope touching market segments in many areas. These include SLA Management, Service Oriented Infrastructures, Cloud computing, Enterprise Service Buses and XaaS provisioning (including Platform as a Service, Software as a Service, Infrastructure as a Service, etc). Through the development of several use cases the project deals with the con-

cerns of a very interesting and diverse cross section of industry, namely ERP hosting, Enterprise IT, Service Aggregation and eGovernment.

#### Outlook

SLA@SOI clearly provides significant progress beyond the current state-of-the-art as it will realize the first service-oriented infrastructure that allows for comprehensive SLA management support across all layers. In such a scenario, the composition of a complex service becomes relatively easy by combining existing services in a dynamic manner. This includes SLA management support that enables a real end-to-end quality of service. Using open interfaces and the combined functionality will also lower the entry barriers for all software providers to participate in a service ecosystem. Thus, small and medium sized enterprises will greatly benefit from SLA@SOI as such companies will be able to deliver very reliable service to their customers.

#### Links:

SLA@SOI: <http://sla-at-soi.eu/>

Project Twitter: <http://twitter.com/slasoi>

#### Please contact:

Wolfgang Theilmann  
SAP Research, CEC Karlsruhe,  
Germany  
Tel: +49 6227 7-52555  
[wolfgang.theilmann@sap.com](mailto:wolfgang.theilmann@sap.com)

Ramin Yahyapour  
IT and Media Center/Service  
Computing Group  
Technische Universität Dortmund,  
Germany  
Tel: +49 231 755-2346  
E-mail: [ramin.yahyapour@udo.edu](mailto:ramin.yahyapour@udo.edu)

# BEinGRID Presage of the Cloud

by Daniel Field

**One of FP6's largest projects recently came to a successful conclusion. Over the last four years the On-Demand IT services sector has transformed beyond recognition, both in commercial and research spheres. Here's how BEinGRID's legacy lives on in today's cloud environment.**

Back in 2005 it was clear that the prevailing Grid technology paradigm was maturing along numerous technical lines. Efforts were seen to incorporate aspects of SOA, multiagent systems, and semantics. We saw nontechnical aspects such as Grid economics being developed and improvements to the core technologies. However, despite the enormous potential of this technology, real case studies seemed scant. When delivering a pitch for a Grid implementation the same sticking point would appear again and again: "it looks and sounds great, but who else in my business has done it successfully?"

It was the BEinGRID Project that strove to resolve this deadlock. This was a research project with an innovative twist: what if instead of developing new technology and testing it a pilot trial, with the inevitable result that it would never be a perfect match for a second implementation, one started with the existing technology and implemented it in multiple pilots, developing only the missing bits and drawing conclusions for future implementations? What if the project was agnostic to the exact technologies applied? What if

what really mattered was the business outcome?

As a consequence, 98 organisations from all across Europe united in what was the FP6's largest Grid project. Grouped into 25 different pilots, each with an end-user, Grid provider and specialist, they focused on the real business problems faced by the end user and built a solution around them. Each was highly autonomous in their solutions: some used GRIA, others Globus, Glite or Unicore. Some were open source, others proprietary. Some were to be delivered as SaaS, others used in-house. Some were highly successful and others, inevitably, less so.

However, the success or failure of any one pilot was arbitrary to the main goals of the project: the project sought to understand the requirements for commercial uptake; to validate the adoption of the technologies by business; and most of all to develop a critical mass of Grid-enabled pilots, across a broad spectrum of economic sectors.

To this end, a core group of organisations led by international IT services

company Atos Origin,, established two groups of consultants, one technical, one business, to work with the pilots, nurturing and advising them during their set up, working with them to develop generic components used across the pilots, as well as ensuring a focus on the business requirements and the steady development of a business plan. As the project progressed the mentors took a step back: what had worked and what hadn't? What were the rules of thumb for applying the technology? Which tendencies were the business-led pilots showing that would impact the future of the research agenda?

Indeed the very trends that were soon to shake up the sector were observed in the pilots from day one: applications had to have simplicity, to be delivered over the network, to be pay-as-you-go, and it had to be affordable. In return the businesses were willing to relinquish some of the tight control they had traditionally enjoyed, indeed in some cases they didn't even want to know what was going on 'under the hood'.

Of course it wasn't long before these characteristics came to define the cloud

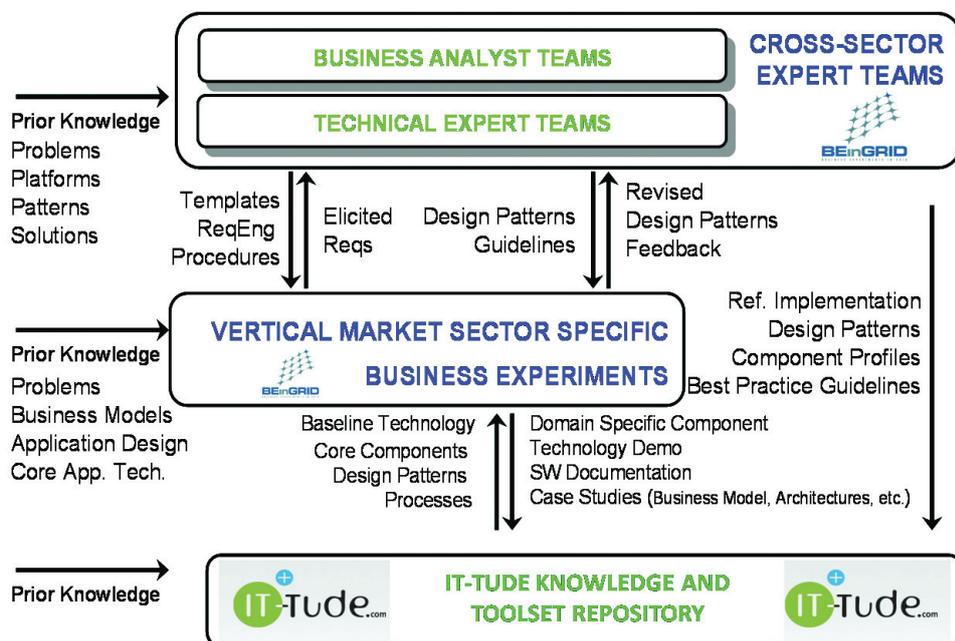


Figure 1: Organisation of the IT-Tude project.

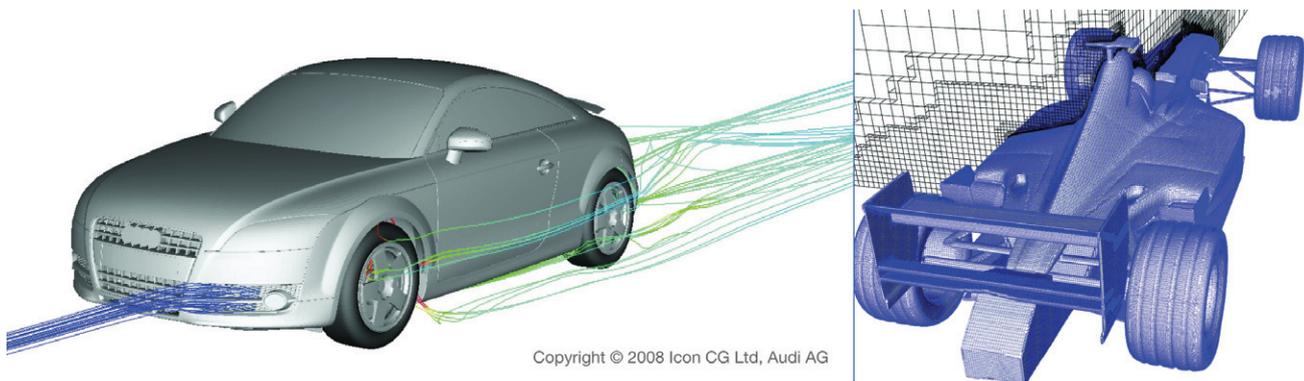


Figure 2: Computational Fluid Dynamics - one of the end user market sectors.

computing phenomenon that swept through the IT industry like a hurricane. By the end, many of the pilots were basing their future plans firmly in the cloud. However the legacy of BEinGRID was not the cloudification of isolated companies, but an entire body of knowledge on the requirements, business drivers, technical hurdles (with solutions), preliminary results and business potential for the migration of a conventional business solution to the cloud computing paradigm.

These results are available to the public in a variety of formats. Three main books have been authored by the project: a short book, "Approaching the Cloud: Better Business Using Grid Solutions" is available for download from the projects website [www.bein-grid.eu](http://www.bein-grid.eu), and the full length books "Service Oriented Infrastructures and Cloud Service Platforms for the Enterprise - A selection of common capabilities validated in real-life business trials by the BEinGRID consortium" and "Grid and Cloud Computing - A Business Perspective on Technology and Applications" are available for purchase from Springer.

Furthermore, BeinGRID collected and published the results, augmenting them with numerous other articles and software from other initiatives through the portal [www.it-tude.com](http://www.it-tude.com). IT-Tude quickly became a reference point for the cloud community. It continues to be developed, supported by the organisations Atos Origin, CETIC, EPCC and NTUA. Highly recommended is the case study library, with close to 100 case studies.

Four professional demonstration kits were developed comprising videos, software, live demonstrations and other material. Proponents of cloud technologies can freely use these to demonstrate to potential clients how the technologies have been applied and the business value that has been derived from them.

The legacy of the project is larger too than just the results gathered whilst the project was active. Many of the partners have gone on to further develop cloud solutions based on their experiences, as is the case for Atos origin, which has launched the commercial cloud Atos Sphere, and numerous lessons have been learned by all the participants. One of these experiences is the management of a large consortium and the success of the open call method to attract top notch end users. Based on the success of this open call system, it is being rolled out across the entire Future Internet Research and Experimentation programme, being developed as part of the Future Internet Assembly. Incidentally, one of these projects, the BonFIRE project, which is coordinated by BEinGRID coordinator Atos Origin, will soon issue an open call for researchers from the cloud computing community who wish to experiment on their generic cloud testbed.

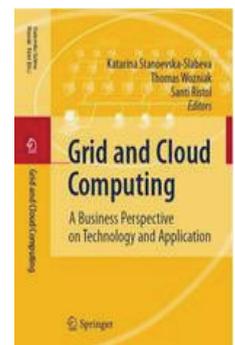
And thus the results, experiences and conclusions of BEinGRID live on. Both software and know-how continues to be applied in the commercial and research activities of its participants, the project's published results are available and used by third parties, and we see the application of intangible experience and working practices to future initiatives.

**Further reading:**

"Service Oriented Infrastructures and Cloud Service Platforms for the Enterprise - A selection of common capabilities validated in real-life business trials by the BEinGRID consortium". Dimitrakos, Theo; Martrat, Josep; Wesner, Stefan (Eds.) 2010, XV, 210 p., Hardcover Springer - ISBN: 978-3-642-04085-6



"Grid and Cloud Computing - A Business Perspective on Technology and Applications" Stanoevska-Slabeva, Katarina; Wozniak, Thomas; Ristol, Santi (Eds.). 2010, X, 274 p., Hardcover Springer - ISBN: 978-3-642-05192-0



**Links:**

<http://www.beingrid.eu>  
<http://www.it-tude.com>  
<http://www.atossphere.com>

**Please contact:**

Daniel Field,  
 Atos Origin  
 E-mail: [Daniel.Field@atosorigin.com](mailto:Daniel.Field@atosorigin.com)

# From XtreamOS Grids to Contrail Clouds

by Christine Morin, Yvon Jégou and Guillaume Pierre

*XtreamOS is an open-source distributed operating system for large scale dynamic Grids. It has been developed in the framework of the XtreamOS European project funded by the European Commission under the FP6.*

XtreamOS can be seen as an alternative to traditional Grid middleware, facilitating the use of federated resources for scientific and business communities. The XtreamOS operating system provides for Grids what a traditional operating system offers for a single computer: abstraction from the hardware and secure resource sharing between different users. When a user runs an application on XtreamOS, the system automatically finds all resources necessary for the execution. It simplifies the user's work by giving them the illusion of using a traditional computer. XtreamOS supports legacy Linux applications as well as Grid-aware MPI and SAGA applications. Applications can be run in the background or interactively. The latter option allows the use of numerical simulation platforms such as Matlab on the Grid. It also considerably eases Grid application debugging.

The XtreamOS system provides three major services to users: application execution management (AEM), data management (XtreamFS) and virtual organization management (X-VOMS). The application execution manager provides scalable resource discovery through a peer-to-peer overlay which connects all self-described resources. XtreamOS provides location-independent access of user data through XtreamFS (<http://www.xtreamfs.org>), a Posix compliant file system spanning the grid. User management in XtreamOS is delegated to virtual organization managers. Access rights to resources are based on policies. Policy rules are defined by virtual organizations as well as by administration domains. They are checked at reservation time and are enforced on resources during execution.

## Cloud Computing: A New Playground for XtreamOS

While XtreamOS was originally designed for Grids, it now appears to be an attractive technology for Cloud computing. During the last two years, we conducted a number of feasibility studies demonstrating that XtreamOS is

highly relevant in the context of virtualized distributed computing infrastructures.

Infrastructure as a Service (IaaS) refers to systems that provide their users with computing resources delivered as a service, including servers, network equipment, memory, CPU and disk space. Although the term was coined after the start of the XtreamOS project, it precisely describes the goal of XtreamOS: to provide users with computing resources that can be assigned to them dynamically when required. In particular, AEM and XtreamFS can legitimately be classified as IaaS components.

AEM allows users to reserve machines through an XtreamOS Virtual Organization (VO) when they need to execute a job. In this sense, it is directly comparable to Amazon's EC2 and other similar services. The two types of services differ in four main aspects:

1. Different computation-as-a-service offers rely on different APIs. So far no clear consensus has emerged regarding a standard access API for such services. On the other hand, AEM is meant to be invoked via the XOSAGA API, which relies on the standard SAGA API for Grid Applications. One could relatively easily build an EC2-compatible API to the AEM.
2. IaaS services typically rely on virtualization techniques where the resources are offered to users in the form of a full virtualized operating system instance. On the other hand, the AEM executes jobs directly, with no virtualization. This difference arises from different requirements in different services. Cloud platforms must be usable by a very large range of users each of whom may want to use a different operating system and work in total isolation from others. On the other hand, XtreamOS intends to be the standard operating system used to

develop Grid applications, so there is no need to virtualize XtreamOS on top of XtreamOS. XtreamOS also provides strong isolation between multiple jobs running on the same hardware through the use of Linux containers.

3. IaaS platforms use a pay-as-you-go pricing model, while XtreamOS relies on the trust relationships between system administrators of a VO to implement a shared resource available to all users of the VO.
4. The security of IaaS platforms relies on a one-to-one trust relationship between the cloud provider and the cloud customer. On the other hand, the VO support in XtreamOS allows one to support several potentially mutually distrustful Cloud providers, allowing Cloud customers to select the provider of their choice.

XtreamFS allows Grid users to store data efficiently and share them across the whole system. In this sense, it is directly comparable to Amazon's S3 and other similar services. Again, no standard API for Cloud storage seems to have emerged yet. Most Cloud storage services provide simplistic functionality, allowing a user to write blocks of files that can be read later but remain immutable. Conversely, XtreamFS implements the full Posix API where files can be updated and overwritten. One could relatively easily build an S3-compatible API to XtreamFS.

Although XtreamOS was not originally designed for Cloud computing applications, it does provide a good base platform for developing advanced Cloud computing functionality. We selected the specific topic of scalable database support to demonstrate how one can deploy Cloud functionality on XtreamOS.

Relational databases such as Oracle have been popular for decades. However, the great expressive power of the SQL query language makes it very

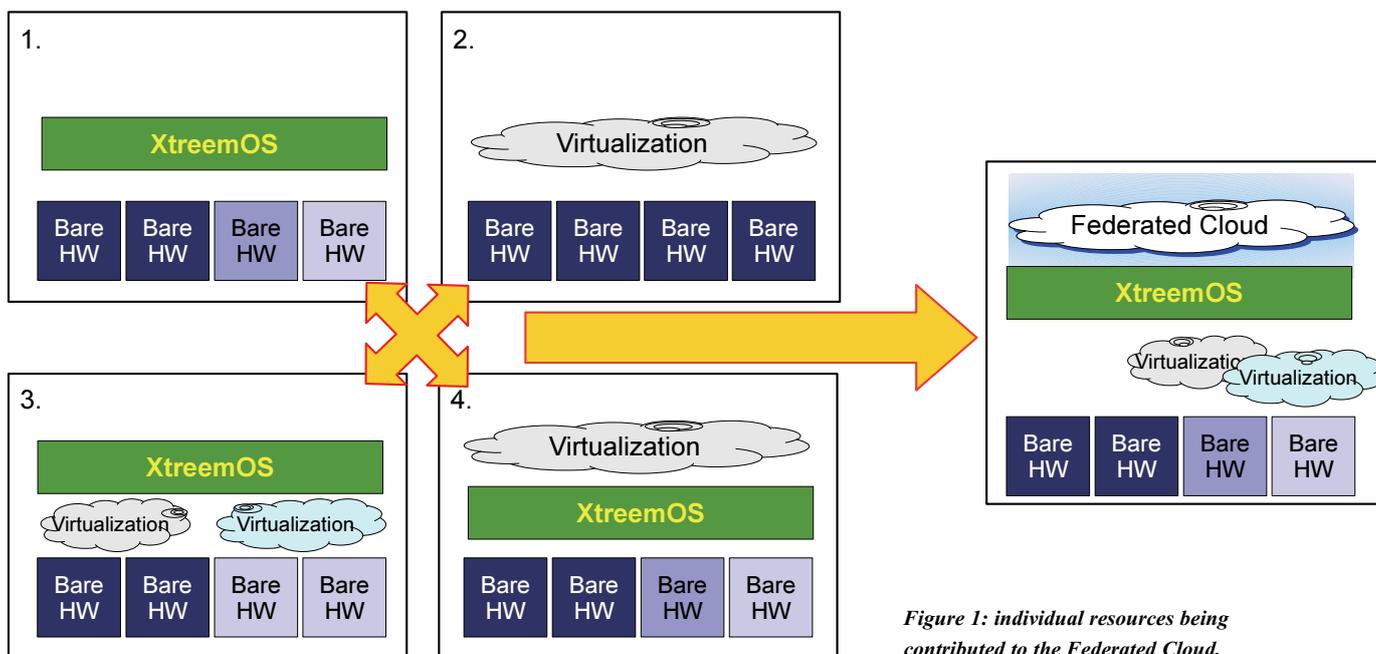


Figure 1: individual resources being contributed to the Federated Cloud.

difficult to scale them up by using large numbers of computers instead of a single powerful database server.

A new family of scalable database systems is being developed for Cloud computing environments, exemplified by Amazon.com's SimpleDB, Google's Bigtable, Yahoo's PNUTS and Facebook's Cassandra. These systems scale nearly linearly with the number of servers they are using, thanks to the systematic use of automatic data partitioning. On the other hand, they do not support the SQL language but rather provide a simpler query language. Data are organized in tables, which can be queried by primary key only. Similarly, these systems do not support join operations. As restrictive as such limitations may look, they do allow construction of useful applications.

To demonstrate how XtreemOS can be a great platform for PaaS Cloud computing, we ported the HBase system (an open-source clone of Bigtable) to XtreemOS. This provides XtreemOS with a scalable database service that can be used by Grid applications to store and query their structured data. Our performance evaluations show that HBase performs well on XtreemOS and allows Grid developers to write scalable data-intensive applications easily.

#### Contrail: an Integrated Approach to Virtualization

The experience acquired in the design of distributed operating systems for the Grid can be exploited in order to deliver

in a timely fashion a system for dependable federated Clouds. The goal of the Contrail project is to develop, evaluate and promote an open source system for Cloud Federations. Contrail will leverage and extend the results from the XtreemOS project. As illustrated in Figure 1, the individual resources being contributed to the Federated Cloud will be highly heterogeneous in their hardware configuration and system-level organization. They may take the form of physical machines running the XtreemOS system (see panel 1), virtual instances from external Clouds (panel 2), virtual machines running XtreemOS (panel 3), or XtreemOS machines running virtualization software (panel 4).

Contrail will vertically integrate an open-source distributed operating system for autonomous resource management in Infrastructure-as-a-Service environments, and high level services and runtime environments as foundations for Platform-as-a-Service. The main achievement will be a tightly integrated software stack in open source including a comprehensive set of system, runtime and high level services providing standardized interfaces for supporting cooperation and resource sharing over Cloud federations.

Contrail will address key technological challenges in existing commercial and academic Clouds: the lack of standardized rich and stable interfaces; limited trust from customers; and relatively poor Quality of Service (QoS) guarantees and SLA support regarding the per-

formance and availability of Cloud resources. Addressing these important issues is fundamental to support large user communities formed of individual citizens and/or organizations relying on Cloud resources for their mission-critical applications.

#### Links:

<http://www.xtreemos.eu>

SimpleDB:

<http://aws.amazon.com/simpledb>

Contrail project:

<http://www.contrail-project.eu>

HBase:

<http://hadoop.apache.org/hbase/>

#### Please contact:

Christine Morin, Yvon Jégou  
INRIA Rennes - Bretagne Atlantique,  
France

E-mail: [christine.morin@inria.fr](mailto:christine.morin@inria.fr),  
[yvon.jegou@inria.fr](mailto:yvon.jegou@inria.fr)

Guillaume Pierre

VU University Amsterdam, The  
Netherlands

E-mail: [gpierre@cs.vu.nl](mailto:gpierre@cs.vu.nl)

# Interoperability between Grids and Clouds

by Attila Marosi, Miklós Kozlovsky and Péter Kacsuk

*Following on from previous successful grid related work, the Laboratory of Parallel and Distributed Systems (LPDS) of SZTAKI is now focusing on Grid-Cloud interoperability.*

Over the last decade the e-science infrastructure eco-system has been enriched with clouds. Together with supercomputer-based grids, cluster-based grids and desktop grids, clouds now form one of the main pillars of this ecosystem. All pillars have their own advantages that make them attractive for a certain application area, however, some of them can benefit from multiple e-science infrastructure technologies. Cluster- and supercomputer-based grids can be considered as ‘service grids’ since they provide managed cluster and supercomputer resources as services with high availability. The recently emerged cloud systems can be used when fast response times, strict Quality of Service (QoS) and Service Level Agreements (SLA) are required and if (as happens in many cases) resource requirements are higher than those provided by the available grid infrastructure resources. However, the cost of porting complex research applications from one technology to another one is hard to finance. Also, unfortunately, in many cases the pillars are separated from each other and

cannot be used simultaneously by the same e-scientist with a large-scale single application.

There are two possible approaches to bind the pillars. The first approach is referred as low-level interoperability which is realized at the task level. Sending tasks directly from one pillar to other one is made possible by using so-called ‘bridging technologies’. The Enabling Desktop Grids for e-Science (EDGeS) FP7 EU project developed a service called 3G-Bridge which allows task based interoperability between service and desktop grids. The recently commenced European Desktop Grid Initiative (EDGI) EU FP7 project continues the work and will develop desktop grid-cloud bridging middleware with the goal to provide instantly available additional resources for desktop grid systems if an application has QoS requirements that cannot be satisfied by the available resources of the system. EDGI will support user communities that are heavy users of e-science infrastructures and require an

extremely large number of CPUs and cores. EDGI is coordinated by the Laboratory of Parallel and Distributed Systems of SZTAKI.

The second approach can be referred as high-level or workflow-level interoperability which allows executing (parts of) workflows on any pillar of the e-science ecosystem. Tasks can be sent to any pillar directly but the low-level interoperability can also be used when required. This type of high interoperability is achieved by using high-level tools like portals (sometimes referred as science gateways). In Europe, one of the most popular generic purpose science gateways is the P-GRADE (Parallel Grid Run-time and Application Development Environment) portal developed by LPDS of SZTAKI. Its basic concepts were mainly developed during the grid era, therefore P-GRADE portal is currently used by many national grids (UK NGS, Belgium Grid, SwissGrid, Turkish/Hungarian/ Spanish Grids, Grid Ireland, etc.), as well as several regional grids (South-East European/Baltic/ UK White Rose Grids) and several science specific virtual organizations (Chemistry Grid, Economy Grid, Math Grid, etc.). In recent years P-GRADE portal has also become popular outside Europe: in Grid Malaysia, Grid Kazakhstan, and Armenian Grid. P-GRADE portal is an open source toolset consisting of a service-rich, workflow oriented graphical front-end and a back-end enabling execution in various types of e-science systems (grids, clusters and recently clouds). It hides the complexity of the infrastructure middleware through its high-level graphical web interface, and it can be used to develop, execute and monitor workflow applications on service and desktop grid systems. P-GRADE portal installations typically provide the user with access to several middleware technologies, using a single login.

Originally, P-GRADE portal supported job submission only to service grids. The 3G-Bridge, developed by EDGeS,

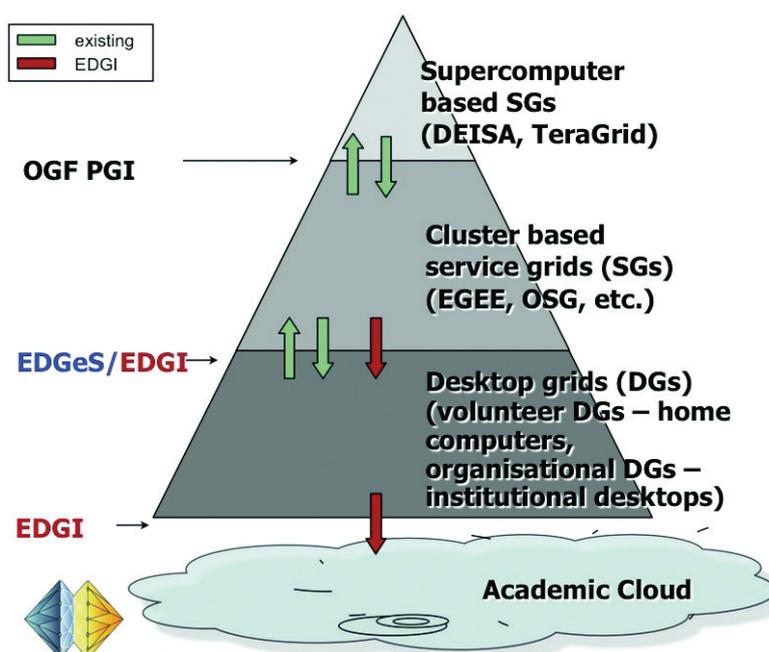


Figure 1: The EDGI project goals.

has been interfaced P-GRADE portal with BOINC and XtremWeb desktop grids, and cloud interoperability is now achieved in three steps:

1. Cloud resource management: allocates on-demand cloud resource for jobs when they arrive and removes the cloud resources when there are no more jobs.
2. Job submission: Submitted jobs are first queued at the 3G-Bridge and then delegated to the allocated cloud resources.
3. Job scheduling on cloud resources: load balances jobs on allocated cloud resources, thus ensures that no

resource is idle while others are overwhelmed.

This approach allows high-level interoperability using on-demand cloud resources in a transparent and efficient way with improved scalability. The three steps (resource management, job submission, job scheduling) of interoperability are separate independent functional units. As an attractive advantage, this allows the independent extension of each component to support different cloud middlewares or user communities with special requirements.

#### Links:

<http://portal.p-grade.hu>  
<http://www.sztaki.hu>  
<http://www.lpds.sztaki.hu>  
<http://edgi-project.eu/>  
<http://edges-grid.eu/>

#### Please contact:

Attila Marosi, Miklós Kozlovsky and Péter Kacsuk  
SZTAKI, Hungary  
E-mail: [atisu@sztaki.hu](mailto:atisu@sztaki.hu),  
[m.kozlovsky@sztaki.hu](mailto:m.kozlovsky@sztaki.hu),  
[kacsuk@sztaki.hu](mailto:kacsuk@sztaki.hu)

## Open Cloud Computing Interface: Open Community Leading Cloud Standards

by Andy Edmonds, Thijs Metsch, Alexander Papaspyrou and Alexis Richardson

*The Open Cloud Computing Interface (OCCI) comprises a set of open community-lead specifications delivered through the Open Grid Forum, which define how infrastructure service providers can deliver their compute, data, and network resource offerings through a standardized interface. OCCI has a set of implementations that act as its proving-ground. It builds upon the fundamentals of the World Wide Web by endorsing the proven REST (Representational State Transfer) approach for interaction and delivers an extensible model for interacting with “as-a-Service” services.*

The aim of the Open Cloud Computing Interface (OCCI) is the rapid development of a clean, open specification and API for cloud offerings. The current focus is on Infrastructure-as-a-Service (IaaS) based offerings but this might extend in future to Platform and Software as a Service offerings. IaaS is one of three primary segments of the emerging cloud computing industry in which compute, storage and network resources are provided as a services. The API is based on a review of existing service provider functionality and a set of use cases (see [1]) contributed by the working group members. OCCI is a boundary API that acts as a service front-end to an IaaS provider’s internal infrastructure management framework. It is OCCI that provides for commonly understood semantics, syntax and a means of management in the domain of consumer-to-provider IaaS. It covers management of the entire life-cycle of OCCI-defined model entities and is compatible with existing work such as the Open Virtualisation Format (OVF). Notably, it serves as an integration point for other standards efforts including DMTF, IETF and SNIA (eg see [2]) as

well as research efforts such as SLA@SOI and RESERVOIR.

OCCI began in March 2009 and was initially lead by co-chairs from the once SUN Microsystems, RabbitMQ and the Universidad Complutense de Madrid. Today, the working group has a membership of over 250 members and includes numerous individuals, industry and academic parties. Some of these members that have contributed include:

- Industry: Rackspace, Oracle, Platform Computing, GoGrid, Cisco, Flexiscale, ElasticHosts, CloudCentral, RabbitMQ, CohesiveFT, Cloud-Central.
- Academia & Research: SLA@SOI, RESERVOIR, Claudia Project, OpenStack, OpenNebula, DGSI.

The reasons driving the development of OCCI were identified as:

- Interoperability - Allow for different Cloud providers to work together without data schema/format translation, facade/proxying between APIs and understanding and/or dependency on multiple APIs

- Portability - No technical/vendor lock-in and enable services to move between providers allows clients easily switch between providers based on business objectives (eg cost) with minimal technical cost and enables and fosters competition.
- Integration: Implementations of the specification can be implemented with those with the latest or legacy infrastructure

Existing specifications for IaaS are provided by single vendors whereas OCCI is the first multi-vendor, community-based initiative to deliver a royalty-free, open standard API. OCCI will improve IaaS interoperability and increase competition. OCCI is also an important enabler for the creation of hybrid cloud architectures that bridge multiple data centers and cloud services. This and related open standards will lower the cost of migration to and from public clouds, delivering automated management of peak capacity, minimizing outages and reducing costs.

An example of how OCCI acts as an enabler for the creation of hybrid cloud

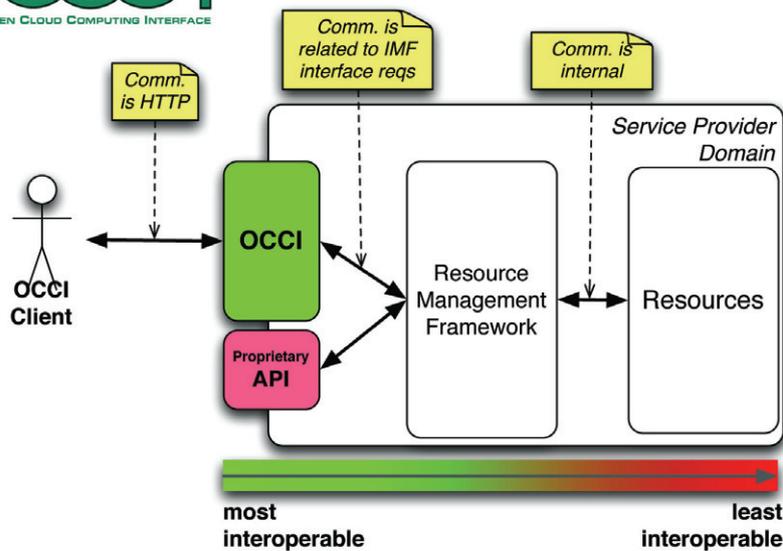


Figure 1: Open Cloud Computing Interface architecture.

architectures is demonstrated through the successful collaboration of two major European Framework Programme 7 (FP7) research integrated projects; SLA@SOI and RESERVOIR. Both projects require the dynamic computing model offered by IaaS and each needed a means to interact with resources managed by IaaS. By choosing to implement the OCCI specification, the collaboration of both projects was significantly eased allowing for the interoperation of two very different infrastructural stacks. The importance of this ease of interoperation cannot be underestimated, especially in the context of large European projects, given their size. Further details of this work can be found in a technical report (see [3]) and a soon to be released white paper.

The OCCI community works in a distributed, open community under the umbrella of the Open Grid Forum (OGF), using a wiki and mailing list for collaboration (see links below). The governance model ensures rights for every voice through the OCCI working group as an open body. Anyone can join free and participate. The OGF's open process is comparable to Standard-bodies such as its sister organization IETF and the complete specification along with any companion documents are publicly and freely available in accordance with OGF's Intellectual Property Rules. The OCCI working group will extend, contribute to and assist the efforts of other groups throughout the process. Having exam-

ined existing APIs as well as developing requirements through the collection of use cases, the group has been able to rapidly produce an implementable technical specification and implementations.

Looking forward, there is continuing work on-going with a number of OCCI implementations. Some interesting implementations include:

- OpenNebula, part of Ubuntu distributions, already supports OCCI,
- SLA@SOI enables automated infrastructure service level agreements using OCCI and
- The Italian National Institute of Nuclear Physics (INFN) are using

OCCI to power their on-demand computing infrastructure

In parallel to this very practical and real work, the specification is currently under revision with excellent work happening in refining the core model of OCCI, seeking opportunities of existing standards reuse as ever and specifying how to present the OCCI core and infrastructural model using semantic web technologies using RDF, RDFa and HTML.

Many of these specification updates and details of implementations will be presented at the upcoming OGF30 conference in Brussels, Belgium.

#### Links:

*Web site:* <http://www.occ-wg.org>

[http://www.ogf.org/Public\\_Comment\\_Docs/Documents/2010-01/occi-http.pdf](http://www.ogf.org/Public_Comment_Docs/Documents/2010-01/occi-http.pdf)

*Wiki:* <http://forge.ogf.org/sf/go/projects.occ-wg/wiki>

*Mailing list:* <http://www.ogf.org/mailman/listinfo/occi-wg>

*SLA@SOI:* <http://www.sla-at-soi.eu>

*RESERVOIR:* <http://www.reservoir-fp7.eu>

*OpenNebula:* <http://www.opennebula.org>

*INFN:* <http://www.infn.it>

*OGF30:* <http://www.ogf.org/OGF30>

[1] OCCI Use Cases: [http://www.ogf.org/Public\\_Comment\\_Docs/Documents/2009-09/occi-usecases.pdf](http://www.ogf.org/Public_Comment_Docs/Documents/2009-09/occi-usecases.pdf)

[2] OCCI and SNIA: <http://www.snia.org/cloud/CloudStorageForCloudComputing.pdf>

[3] Using Cloud Standards for Interoperability Frameworks: <http://sla-at-soi.eu/wp-content/uploads/2010/04/RESERVOIR-SLA@SOI-interop-techReport.pdf>

#### Please contact:

Andy Edmonds, Intel Ireland Limited, Ireland

Tel: +353 87 3074442

E-mail: [andrewx.edmonds@intel.com](mailto:andrewx.edmonds@intel.com)

# Recent Developments in DIET: From Grid to Cloud

by Frédéric Desprez, Luis Roderer-Merino, Eddy Caron and Adrian Muresan

*The Distributed Interactive Engineering Toolkit, or DIET, project started with the goal of implementing distributed scheduling strategies on compute Grids. In recent times, the Cloud phenomenon has you-go billing approach. This led to a natural step forward in the evolution of DIET, with the inclusion of Cloud platforms in resource provisioning. DIET will be used to test resource provisioning heuristics and to port new applications that mix grids and Clouds.*

In 2000, the Grids and Algorithms, or GRAAL, research team of INRIA, located in Ecole Normale Supérieures de Lyon, France, initiated the Distributed Interactive Engineering Toolbox project under the supervision of Frédéric Desprez and Eddy Caron. The project is focused on the development of scalable middleware with initial efforts on distributing the scheduling problem across a hierarchy of agents, at the top of which sits the Master Agent, or MA. At the bottom level of a DIET hierarchy one can find the Service Daemon, or SeD, agents. SeDs are connected to the MA by means of Local Agents, or LAs. Over the last few years, the Cloud phenomenon has been gaining more and more traction in the industry and in research communities because of its qualities, the most interesting of which is its on-demand resource provisioning model

and its pay-as-you-go billing approach. We deem these features to be highly interesting for DIET.

## From Grid to Cloud

The first step towards the Cloud was to enable DIET to take advantage of on-demand resources. This should be done at the platform level and be transparent to the DIET user. The authors and David Loureiro have targeted the Eucalyptus Cloud platform as it implements the same management interface as Amazon EC2, but unlike the latter it allows for customized deployments on the user's hardware. The team has implemented the scenario in which the DIET platform sits completely outside of the Eucalyptus platform and treats the Cloud only as provider of compute resources when needed. This opened the path towards new researches

around grids and Cloud resource provisioning, data management over these platforms and hybrid scheduling in general.

## Cloud application resource scaling

The on-demand provisioning model for resource allocation and the pay-as-you-go billing approach that Cloud systems offer makes possible the creation of more cost-effective approaches for application resource provisioning. A Cloud application can scale its resources up or down to better match its usage and to reduce the number of unused, yet paid for, resources. This leads to smart auto-scaling strategies.

By taking into account research done around self-similarities in web traffic, the authors have developed a resource usage prediction model that identifies similar past resource usage patterns from a historic archive. Once identified, they provide an insight into what the short-term usage of the platform will be. This approach can be used to predict the usage of the most important types of resources of a Cloud client and thus give an insight of what type of virtual machine to instantiate or terminate when the Cloud application is rescaled. We have tested this approach against resource usage traces from one Cloud client and three production grids and obtained encouraging results.

## Economy-based resource allocation

The dynamics that Cloud systems bring in combination with the agent-based DIET platform led us towards an economic model for resource provisioning. The ultimate goal is to guarantee resource sharing fairness and avoid starvation. The authors have done this by simulating the dynamics of a tender/contract-net market. In this market contracts are established between platform users (the DIET clients) and resource providers (the

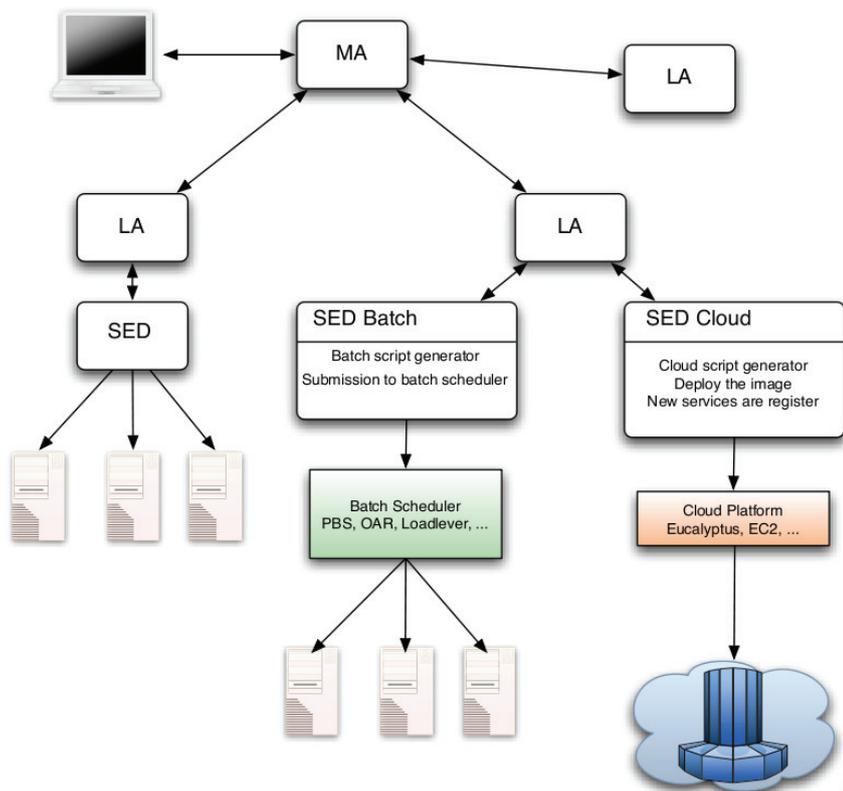


Figure 1: The cloud-enabled DIET hierarchy.

DIET SeDs). Users send requests to the DIET platform for the execution of their tasks and resource providers reply with offers, each containing the cost and duration of the task execution. A user-defined utility function is applied to identify the best offer and the corresponding SeD will run the task.

In this scenario, platform users compete against each other for resource usage while the resource providers compete against each other for profit. Resource prices, which determine the offer costs, fluctuate depending on each provider's resource usage level. Hence, users will

tend to choose SeDs with more free resources and so lower prices.

**What's next?**

Future directions include implementing a complete automatic resource scaling strategy for Cloud clients and testing against real-life situations.

We are also looking forward towards integrating Cloud-specific elements into the DIET scheduler for existing applications. The final goal is to see if deployment on a Cloud platform would yield a better performance and if so then with what scheduling modifications.

Finally, we plan to study hybrid scheduling strategies mixing static grids and dynamic Clouds for a more efficient resource management of large scale platforms.

**Link:**

The DIET project:  
<http://graal.ens-lyon.fr/DIET>

**Please contact:**

Adrian Muresan  
Ecole Normale Supérieure de Lyon,  
France  
Tel: +33 4 37 28 76 43  
E-mail: [Adrian.Muresan@ens-lyon.fr](mailto:Adrian.Muresan@ens-lyon.fr)

## Addressing Aggregation of Utility Metering by using Cloud – The Power Grid Case Study

by Orlando Cassano and Stéphane Mouton

*Utility grids are generating increasingly huge amounts of metering information. Grid operators face rising costs and technical hurdles to aggregate and process data. Can Cloud Computing tools, developed notably by Web companies to deal with large data sets, also be used for power grid management?*

As electricity in the current state of technology cannot be stored, consumption in power grids is continuously counterbalanced by production. Both producers and consumers are connected to the grid through metered Access Points (AP). Every month, grid operators have to determine the amount of energy produced or used by each stakeholder, knowing that the sum of produced energy, either out of the grid by power

plants or within the grid (eg by wind turbines), equals the sum of consumed energy, from effective use and losses in the grid. Amounts are aggregated in order to obtain amount of "allocated" energy by stakeholder. The volume of data at stake in the allocation computation depends on the size of the grid. Data are produced every 15 minutes, and datasets may be huge. For example, there are roughly 8 million APs in

Belgium, each producing 96 metering data outputs per day: allocation for a month would therefore require handling of more than 23 billion records.

Data aggregation is currently based on existing Relational DataBase Management Systems (RDBMS). However performance of such software is declining with the increasing volume of data to process. Performance can be improved by investing in hardware and sophisticated software setups, like database clusters, but such an investment is not necessarily economical, with the cost of such a setup increasing disproportionately in relation to data processing capacity.

The goal of our research was to overcome the limitations of RDBMS by scaling performance according to growth of aggregated data. Moreover the allocation algorithm is a good candidate for parallelization as sums have to be performed on distinct data sets, ie, per stakeholder. For this reason we investigated the use of programming platforms and frameworks, identified as providing Platform as a Service (PaaS) on Cloud infrastructures, to enable scalability in data storage and processing.

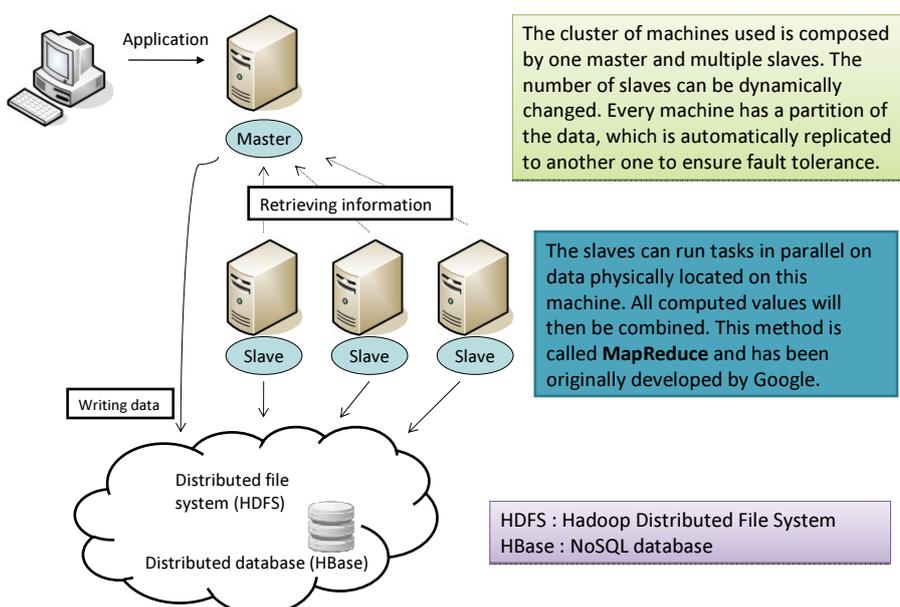


Figure 1: Set up of the Cloud architecture.

Data stored in Cloud platforms can be structured in non relational schema following an approach known as NoSQL – standing for Not Only SQL. Major players on the Web use NoSQL databases to deal with large amounts of data such as Google (Bigtable), Amazon (SimpleDB), Facebook (Cassandra), LinkedIn (Voldemort), etc. NoSQL databases follow several approaches:

- Key/Value oriented databases aim to be a simple abstraction for a file system, acting like a hash table.
- Document-oriented databases extend the key/value database concept to add more information to a key (think object oriented).
- Column-oriented databases associate a set of column families with a key. Each column family, containing a variable set of columns to provide a flexible structure for data, will be stored in different machines.

We found that column-oriented databases would fit our needs best. The storage of data by columns is well designed to make aggregations of the same information for all records in the database.

We used the Open Source Hadoop platform for our implementation. Hadoop

includes a distributed file system (HDFS) – files are “split” over multiple machines – together with a column-oriented database (HBase).

A parallel algorithm has been written to solve the problem of measurements accounting by aggregating metering information located on slaves. Existing collected data have been ported from the relational database to the NoSQL database. The data have been restructured to fit the column-oriented structure. This implementation provides scalability as well as reliability of the data. Indeed, in the set up cluster, more (cheap) machines can be added dynamically to ensure scalability instead of investing in a new (expensive) big server. In addition, automated redundancy of data, provided by many Cloud platforms (in HDFS in our case), improves reliability.

Our first results show, however, that performances of our version of the algorithm ported on Cloud platform are lower than the original implementation. We traced the origin of the problem to data structure. Parallelizing processing is not enough: NoSQL databases require full data reorganization. Data structures used in our first implementation are still

too oriented towards use of RDBMS and we are currently reworking them and adapting our implementation accordingly. This exercise has enabled us to gain better understanding of the core of the metering aggregation in power grid so as to tailor the corresponding solution more effectively.

We remain confident that we will achieve a speedup by using a Cloud platform, due to the parallel nature of the processing. The problem we faced, however, does raise an issue worthy of consideration: if the required skills or effort necessary to implement NoSQL databases is too high, adoption of this new paradigm could be hampered, and the use of such systems might be restrained to new developments without pre-existing use of RDBMS or extreme situations faced by Google or Facebook.

**Link:**

<http://www.cetic.be/article1079.html>

**Please contact:**

Stéphane Mouton  
SST-SOA team (SOA & Grid Computing)  
CETIC, Belgium  
Tel: +32 71 490 726  
E-mail: [stephane.mouton@cetic.be](mailto:stephane.mouton@cetic.be)

## Optimization and Service Deployment in Private and Public Clouds

by Máté J. Csorba and Poul E. Heegaard

*Large-scale computing platforms will soon become a pervasive technology available to companies of all sizes. They will serve thousands, or even millions of users through the Internet. However, existing technologies are based on a hierarchically managed approach that does not possess the required scaling properties. Moreover, existing systems are not equipped to handle the dynamism caused by severe failures or load surges. We conjecture that using self-organizing techniques for system (re)configuration can improve both the scalability properties of such systems as well as their ability to tolerate variations in load and increased failure rates. Specifically, we focus on the deployment of virtual machine images onto physical machines that reside in different parts of the network. Our objective is to construct balanced and dependable deployment configurations that are resilient and support elasticity. To accomplish this, a method based on a variant of Ant Colony Optimization is used to find efficient deployment mappings for a large number of replicated virtual machine images that are deployed concurrently. The method is completely decentralized; ants communicate indirectly through pheromone tables located in the nodes.*

Central to our work – conducted at the Department of Telematics at NTNU – are distributed software services hosted in hybrid cloud-like environments possibly with multiple providers and their non-functional requirements, for such as those related to system performance and

dependability. We find optimal deployment mappings involving multiple services, ie map service components – VMs in an IaaS scenario – in the software architecture to the underlying platforms for best possible execution. Requirements are used to construct

appropriate cost functions that guide our heuristic optimization method. In particular, we obtain a decentralized method using swarm intelligence, free of discrepancies of centralized solutions, such as single point of failures and performance bottlenecks.

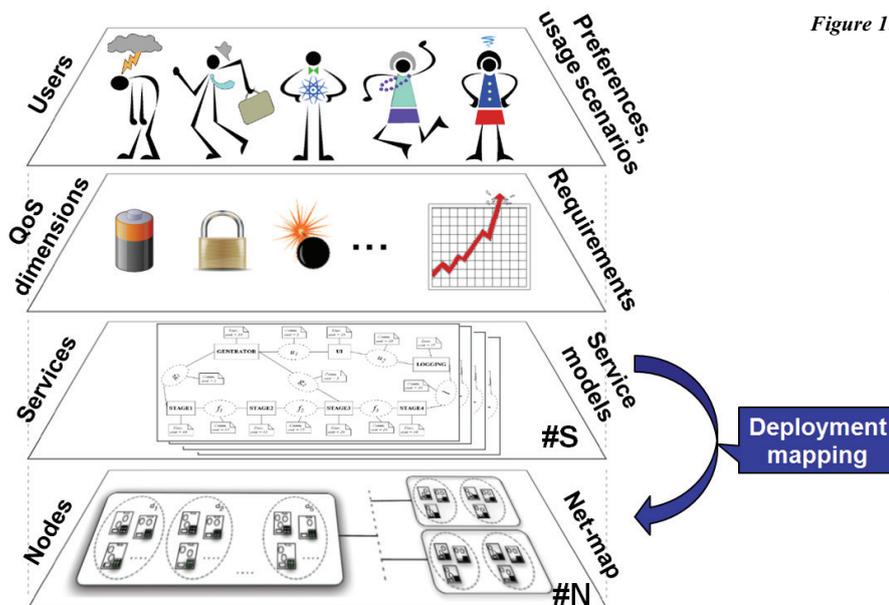


Figure 1: The deployment problem.

Nodes hosting a service may be heterogeneous and may provide a dynamic environment. For example, nodes can join and leave the network in an unpredictable manner, in particular, in large-scale data-centers designed to handle failures, which are present most of the time. Changing context of distributed services requires the capability of adaptation to satisfy QoS-requirements, while taking into account costs from the service providers' perspective. The wide range of possible requirements makes the deployment problem a multi-faceted challenge demanding multi-dimensional optimization (see. Figure 1).

**Making Placement Decisions in a Cloud Computing Setting**

We consider a large-scale data-center consisting of a collection of nodes (see. Figure 2), which can be organized into sets of clusters. The data-center hosts a set of services ( $S_1, S_2, \dots$  in Figure 2 (a)), in which each component ( $V_1^1, V_2^1, \dots$ ) may be replicated ( $R_{11}^1, R_{21}^1, \dots$ ) for fault tolerance and/or load-balancing purposes. The method is implemented in the form of ant-like agents moving in the network to identify potential locations for placement (represented by the green and blue ants in Figure 2 (b)). Different ant species are

responsible for different services. The execution environment has to install, run and migrate VMs. Our method, however, is transparent regarding the execution framework. Nodes have pheromone tables manipulated by visiting ants to reflect their knowledge of possible mappings. These tables are used by ants to select suitable deployment mappings. To find mappings satisfying the requirements, the Cross-Entropy Ant System (presented in ERCIM News No. 64) is used, which works by evaluating the findings using a cost function. CEAS then uses the Cross-Entropy stochastic optimization

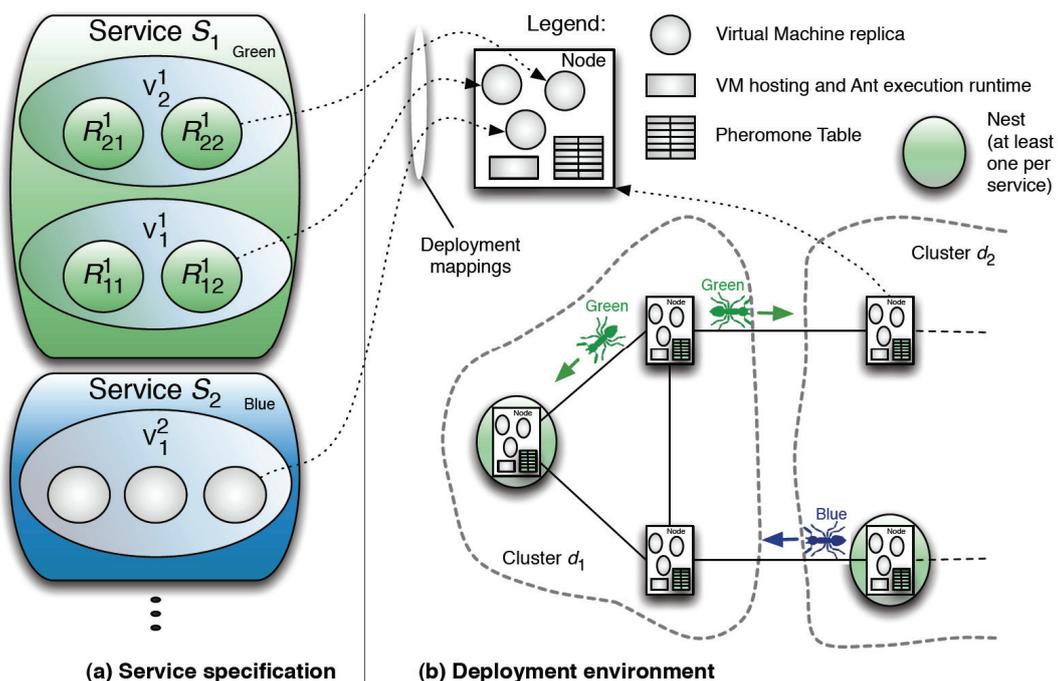


Figure 2: CE Ant System for Deployment Mapping.

method to gradually alter the pheromones according to the cost of the mapping found. To deploy a service, at least one node must be running a nest for that service. Tasks of a nest are: (i) to emit ants for the associated service, and (ii) trigger placement, once a convergence criteria is satisfied.

To demonstrate feasibility of our approach, we have modelled several scenarios and conducted simulations, including tailored examples of traditional NP-hard task assignment problems, deployment problems of collaborating software components within multiple parallel services and deployment of VMs in virtualized computing clouds. In the test scenarios we have looked at the convergence properties of our algorithm, different possible encodings of pheromone values, and obtained

mappings that with high confidence satisfy the requirements of the services provisioned, such as dependability (cluster-disjointness and co-location avoidance) and load-balancing among the nodes. Cross-validation of the results obtained using our distributed logic against centralized solutions for finding mappings has been conducted applying integer linear programming further increasing confidence in our approach.

#### Discussion

We are developing a distributed optimization algorithm based on an ant colony system, the CEAS. It finds efficient mappings between service components, such as VMs in a cloud computing setting and hosts suitable for execution. Identified mappings adhere to predefined non-functional require-

ments specified in formal models. The algorithm is at present implemented in a discrete event simulator environment and we are currently working on porting it to a Java based environment, which would, in the long term, ease experimentation in an online network. Future plans include extending our method with energy-saving aspects that have become key in production data-centers, and further improving scalability and adaptation capabilities of our algorithm.

#### Links:

<http://www.item.ntnu.no/~csorba/>  
<http://www.item.ntnu.no/~poulh/>

#### Please contact:

Máté J. Csorba or Poul E. Heegaard,  
Dept. of Telematics, NTNU, Norway  
Tel.: +47 73590786  
E-mail: {csorba, poulh}@item.ntnu.no

## Holistic Management for a more Energy-Efficient Cloud Computing

by Eduard Ayguadé and Jordi Torres

*For a more sustainable Cloud Computing scenario the paradigm must shift from “time to solution” to “kWh to the solution”. This requires a holistic approach to the cloud computing stack in which each level cooperates with the other levels through a vertical dialogue.*

Due to the escalating price of power, energy-related costs have become a major economic factor for Cloud Computing infrastructures. Our research community is therefore being challenged to rethink resource management strategies, adding energy efficiency to a list of critical operating parameters that already includes service performance and reliability.

Current workloads are heterogeneous, including not only CPU-intensive jobs, but also streaming, transactional data-intensive, and other types of jobs. These jobs are currently performed using hardware that includes heterogeneous clusters of hybrid hardware (with different types of chips, accelerators, GPUs, etc.). In addition to the goal of improved performance, the research goals that will direct research proposals at the Barcelona Supercomputing Center (BSC) include fulfilling the Service Level Agreements (SLA), considering energy consumption and taking into account the new wave of popular pro-

gramming models like MapReduce. These cloud goals however, have made resource management a burning issue in today's systems. For BSC, self-management is considered the solution to this complexity and a way to increase the adaptability of the execution environment to the dynamic behaviour of Cloud Computing. We are considering a whole control cycle with a holistic approach which involves each level cooperating with other levels through a vertical dialogue. Figure 1 shows a diagram that summarizes the role of each approach and integrates our current proposals.

#### Virtual Infrastructure

At Infrastructure-as-a-Service (IaaS) level BSC is contributing the EMOTIVE framework to the research community. This framework simplifies the development of new middleware services for the Cloud. The EMOTIVE framework is an open-source software infrastructure for implementing Cloud computing solutions that provides

elastic and fully customized virtual environments in which to execute Cloud services. One of the main distinguishing features of EMOTIVE framework is its functionalities that ease the development of new resource management proposals, thus contributing to innovation in this research area. Recent work extends the framework with plugins for third-party providers and federation support for simultaneous access to several clouds that can take into consideration energy-aware parameters.

#### Distributed Management

At Platform-as-a-Service (PaaS) level we are working on application placement to decide where applications run and the allocated resources required. To this end, applications must be designed to make proper placement decisions in order to obtain a solution that considers energy constraints as well as performance parameters. We are paying particular attention to MapReduce workloads (currently the most prominent emerging model for cloud scenario), working on

the runtimes that allow control and dynamic adjustment of the execution of applications of this type with energy awareness. Finally the energy awareness is addressed at two levels: compute infrastructure (data placement and resource allocation) and network infrastructures (improving data locality and placement to reduce network utilization).

#### High Level Management

We are considering extending the Platform-as-a-Service layer functions to provide better support to Software-as-a-Service (SaaS) layer, according to high level parameters for resource allocation process. The main goal is to propose a new resource management aimed to fulfil the Business Level Objectives (BLO) of both the provider and its customers in a large-scale distributed system. We have preliminary results that describe the way decision-making processes are performed in relation to several factors in a synergistic way depending on provider's interests, including business-level parameters such as risk, trust, and energy.

#### Exploiting Emerging Hardware

BSC is interested in studying and development of both new hardware architectures that deliver best performance/energy ratios, and new approaches to exploit such architectures. Both lines of research are complementary and will aim to improve the efficiency of hardware platforms at a low level. Preliminary results demon-

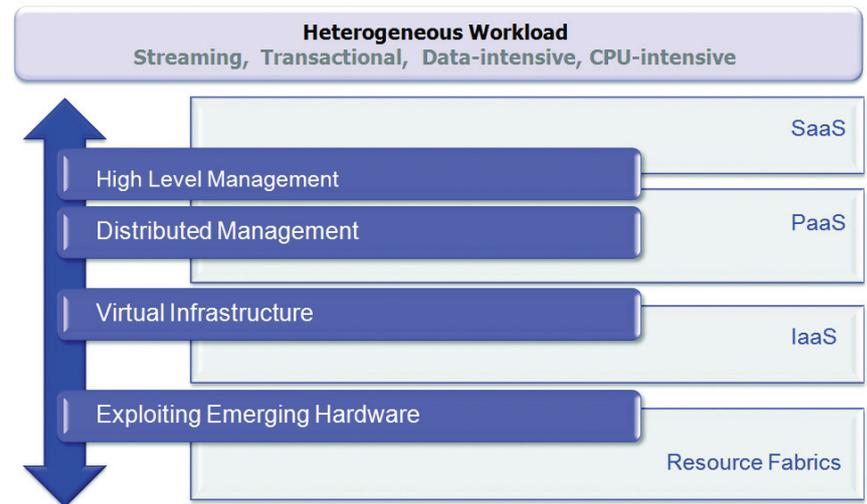


Figure 1: Cloud computing stack organization and BSC contributions.

strate that the energy modelling in real time (based on processor characterization) will be leveraged to make decisions. We are focusing on leveraging hybrid systems to improve energy-saving, thereby addressing the problem of low-level programmability of such systems that can result in poor resource utilization and, in turn, poor energy efficiency.

The Autonomic Systems and e-Business Platforms department at the Barcelona Supercomputing Center (BSC) is proposing a holistic approach to the cloud computing stack in which each level (SaaS, PaaS, IaaS) cooperates with the other levels through a vertical dialogue, trying to build a “Smart

Cloud” that can address the present challenges of the Cloud. The current research at BSC is about autonomic resource allocation and heterogeneous workload management with performance and energy-efficiency goals for Internet-scale virtualized data centres comprising heterogeneous clusters of hybrid hardware.

#### Link:

<http://www.bsc.es/autonomic>

#### Please contact:

Jordi Torres  
Barcelona Supercomputing Center  
UPC Barcelona Tech / SpaRCIM, Spain  
Tel: +34 93 401 7223  
E-mail: [torres@ac.upc.edu](mailto:torres@ac.upc.edu)  
<http://people.ac.upc.edu/torres>

## A Semantic Toolkit for Scheduling in Cloud and Grid Platforms

by András Micsik, Jorge Ejarque, Rosa M. Badia

*Delivering a good quality of service is crucial for service providers in cloud computing. Planning the schedule of resource allocations and adapting the schedule to unforeseen events are the primary means of obtaining this goal. Within the EU funded IST project BREIN (Business objective driven reliable and intelligent Grids for real business), semantic and agent technologies have been applied to implement a platform with scheduling, monitoring and adaptation to ensure the agreed quality of service during service provision. In the Department of Distributed Systems of SZTAKI and the Barcelona Supercomputing Centre novel semantic techniques applied in the platform have been developed, namely prediction of quality of service based on historical data and allocation of licenses.*

From a business point of view, the service provider establishes an agreement with its customers regarding the Quality of Service (QoS) and level of service through a Service Level

Agreement (SLA). The fulfilments or violations of the SLAs indicate the level of customer satisfaction with the Service Provider (SP), affecting directly or indirectly the benefit of these

providers. In the cloud environment a service provider can outsource the resources used to execute services to the public cloud. During such outsourcing the allocation mechanism has to cope

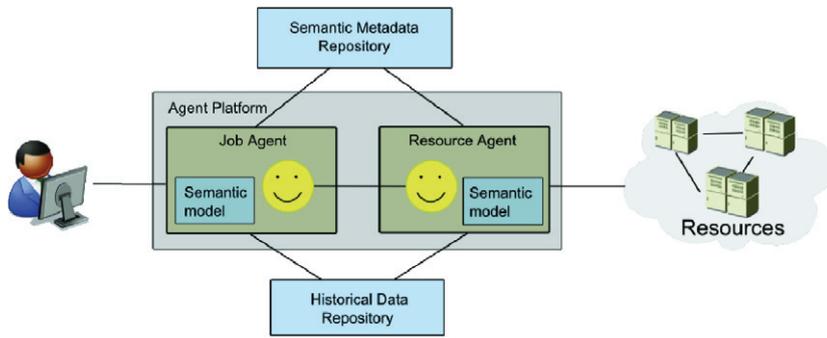


Figure 1: Core architecture for semantic resource allocation.

with the big number of infrastructure providers with their different resource descriptions and the complexity of allocating services to different providers capable of fulfilling the customers' requirements.

The platform, developed in BREIN for semantic resource allocation, applies a multi-agent system (JADE) to distribute the decisions on resource allocation. The multi-agent technology helps coordinate resource allocation and assists adaptation of service execution, while the semantic web technology helps to leverage interoperability with infrastructure providers.

There are two types of agent in the architecture: Job Agents (JA) are in charge of managing the customers' executions; and Resource Agents (RA) are in charge of managing the providers' resources. The scheduling of the jobs in the different resources is made by an agreement obtained from a negotiation between a Job Agent and different Resource Agents. As part of the negotiation, RAs propose allocations according to their provider and the JA selects the allocation proposal most suitable for the customer. During this process agents are supported by two general semantic services: a Semantic Metadata Repository (SMR) containing the cur-

rent semantic resource descriptions registered in the platform, and the Historical Data Repository (HDR) containing semantically annotated logs from system events such as job executions, failures and other monitoring data.

In order to learn from past experiences, agents use the HDR as a service. On job completion, the Job Agent sends the job execution details to the HDR component. Resource Agents report failures to the HDR. All this information is merged in a single semantic store inside the HDR, enabling the generation of statistics and predictions, created on request by agents during the planning or adaptation of resource allocations. Predictors can be installed as plug-ins for HDR to answer specific questions. Within HDR, the RDF (Resource Description Framework) storage is coupled with data mining software. In this way, a predictor can use both semantic querying and statistical methods. For example, a predictor can build a classification model on top of the results of a semantic query. In real use cases HDR was used to predict delays and failures of job executions, and to assess the reliability of hosts.

Another novel feature of the semantic allocation process is the capability to

allocate software licenses for customer requests. The lack of software licenses can block customers' jobs just as jobs may be blocked by a lack of computing resources. However, software license allocation is regulated by quite different rules than computing resource allocation. There exists a huge variety of license restrictions, and thus the machine understandable representation of licenses is a huge and difficult task. In order to provide a practical solution for frequent use cases, we limited the semantic descriptions to the viewpoint of resource management, and compiled an extensible core for the semantic description of software licenses, where issues raised by new types of license can be covered by new rules plugged into the running environment. With this approach we successfully modelled license term restrictions on CPU numbers, temporal limitations, user limitations, and hosts running the software. During license allocation, the set of suitable licenses for a client request is created and filtered by a customizable rule set within the Jena Semantic Web toolkit, where finally only the applicable licenses are sent to the Job Agent for final selection and allocation.

The BREIN project team extended and connected a set of ontologies in order to support the mentioned functionalities. Based on OWL-S and the Grid Resource Ontology (GRO) we created an OWL-DL environment in which business and technical aspects of service provisioning can be described and related to each other.

BSC and SZTAKI have jointly developed the two new features detailed here, and plan to continue to test more real-life use cases and to further extend the use of semantic techniques in adaptable resource allocation and job scheduling.

**Link:**  
<http://www.eu-brein.com/>

**Please contact:**  
 András Micsik  
 SZTAKI, Hungary  
 Tel: +36 1 279 6248  
 E-mail: micsik@sztaki.hu

Jorge Ejarque - Barcelona  
 Supercomputing Center, Spain,  
 Tel: +34 934137248  
 E-mail: jorge.ejarque@bsc.es

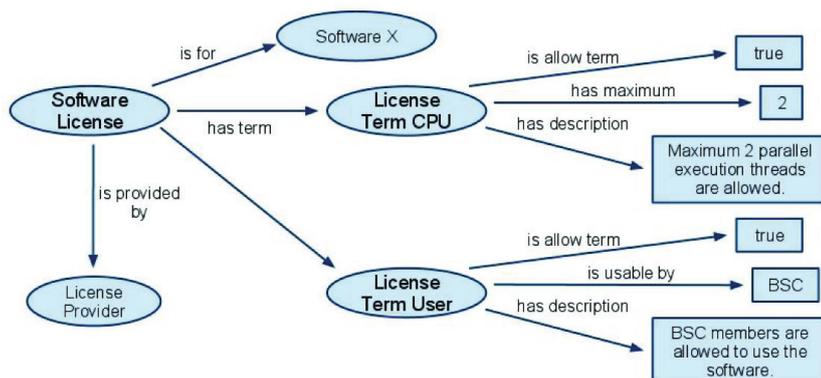


Figure 2: Example for a software license description.

# Making Virtual Research Environments in the Cloud a Reality: the gCube Approach

by Leonardo Candela, Donatella Castelli, Pasquale Pagano

*In recent years scientists have been rethinking research workflows in favour of innovative paradigms to support multidisciplinary, computationally-heavy and data-intensive collaborative activities. In this context, e-Infrastructures can play a crucial role in supporting not only data capture and curation but also data analysis and visualization. Their implementation demands seamless and on-demand access to computational, content, and application services such as those typified by the Grid and Cloud Computing paradigms. gCube is a software framework designed to build e-Infrastructures supporting Virtual Research Environments, ie on-demand research environments conceived to realise the new science paradigms.*

The eScience community is currently examining the feasibility of setting up innovative Virtual Research Environments (VREs) to meet the requirements of collaborative activities. VREs are designed to support both small and large-scale computationally-intensive, data-intensive and collaboration-intensive tasks, and to serve research communities potentially distributed over multiple domains and institutions.

A promising approach for the building and operation of VREs is based on e-Infrastructures, ie frameworks that enable secure, cost-effective and on-demand resource sharing across organizational boundaries. An e-Infrastructure can be seen as a “mediator”, accommodating resource sharing among resource providers and consumers, either human or inanimate. Resources are intended as generic entities, either physical (eg storage and computing resources) or digital (eg software, processes, data), that can be shared and can interact with other resources to synergistically provide various types of service. A service-based paradigm is needed in order to share/reuse these resources. The e-Infrastructure layer allows resource

providers to “sell” their resources, and resource consumers to “buy” them and to use them to build their applications. It also provides organizations with logistic and technical support for application building, maintenance, and monitoring.

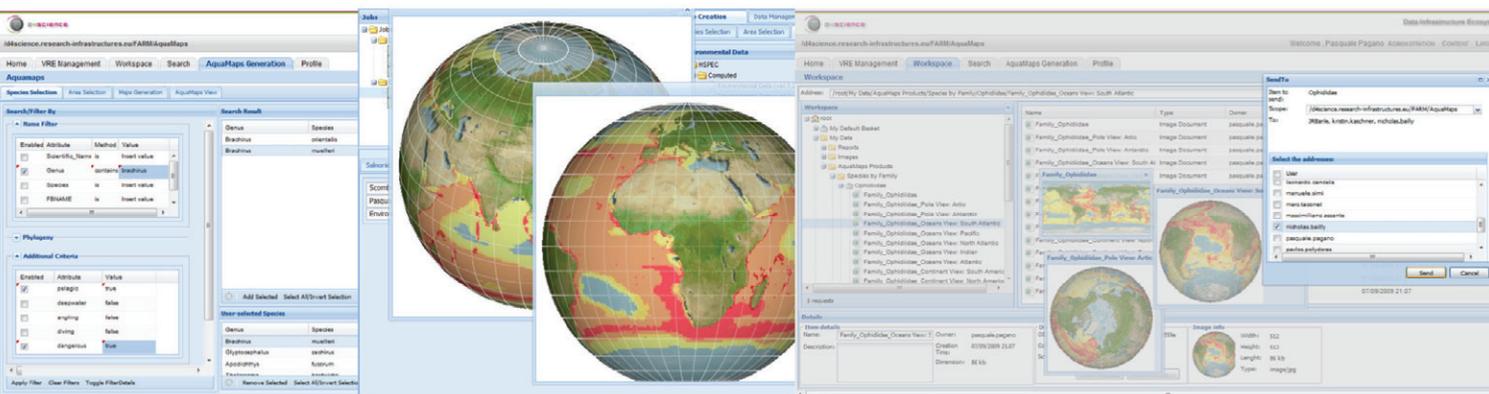
The e-Infrastructure vision shares many commonalities with Grid Computing and Cloud Computing. All three aim to reduce computing costs via economies of scale. They all attempt to achieve this objective by managing a pool of abstracted and virtualized resources and offering on demand computing power, storage facilities and services to “external” customers over the internet. The differences mainly reside in the services they offer, the business models, and the technologies that characterize them.

gCube is a software system specifically conceived to develop and operate large scale e-Infrastructures, enabling the declarative definition and automatic deployment and operation of VREs.

gCube facilities for e-Infrastructure development include a rich array of

mediator services for interfacing with existing infrastructure enabling technologies including grid (eg gLite/EGEE), cloud (eg Hadoop) and data source (eg OAI-PMH) oriented approaches. Via these mediator services, the storage facilities, processing facilities and data resources of the external infrastructures are conceptually unified to become gCube resources. Facilities for deploying gCube Nodes, ie servers offering storage and computing facilities (similar to the Infrastructure as a Service Cloud) are also offered together with the dynamic deployment of gCube services (similar to the Platform as a Service Cloud). These resources are complemented by the Software as a Service Cloud approach, ie offering software frameworks for data management, data integration, workflow definition and execution, information retrieval, and user interface building.

By relying on this impressive amount of resources and services, gCube based e-Infrastructures enable scientists to declaratively and dynamically build the VREs they need while abstracting on the implementation details. gCube tech-



Screenshots of gCube based Virtual Research Environments.

nology implements a user friendly Platform as a Service Cloud “function” where the content, application services, and computing resources needed by a scientist are automatically aggregated and deployed, and made available through a web based interface. The aggregated resources are also monitored to guarantee the VRE service.

gCube technology is now serving a number of challenging scientific domains, for example marine biologists generating model-based large-scale predictions of natural occurrences of marine species, High Energy Physicists mining bibliometric data and producing

hybrid metrics on the entire corpus of their literature, and fishery statisticians managing and integrating catch statistics.

gCube is the result of the collaborative efforts of researchers and developers from academic and industrial research centres including the Institute of Information Science and Technologies ISTI-CNR (IT), University of Athens (GR), University of Basel (CH), Engineering Ingegneria Informatica SpA (IT), University of Strathclyde (UK), CERN European Organization for Nuclear Research (CH), 4D SOFT Software Development Ltd (HU). Its

development has been partially supported by the DILIGENT project (FP6-2003-IST-2, Contract No. 004260), the D4Science project (FP7-INFRA-2007-1.2.2, Contract No. 212488), and the D4Science-II project (FP7-INFRA-2008-1.2.2, Contract No. 239019)

**Link:**

gCube: <http://www.gcube-system.org>

**Please contact:**

Pasquale Pagano

ISTI-CNR, Italy.

E-mail: [pasquale.pagano@isti.cnr.it](mailto:pasquale.pagano@isti.cnr.it)

## ManuCloud: The Next-Generation Manufacturing as a Service Environment

by Matthias Meier, Joachim Seidelmann and István Mezgár

*The objective of the ManuCloud project is the development of a service-oriented IT environment as a basis for the next level of manufacturing networks by enabling production-related inter-enterprise integration down to shop floor level. Industrial relevance is guaranteed by involving industrial partners from the photovoltaic, organic lighting and automotive supply industries.*

The transition from mass production to personalized, customer-oriented and eco-efficient manufacturing is considered to be a promising approach to improve and secure the future competitiveness of the European manufacturing industries, which constitute an important pillar of European prosperity. One precondition for this transition is the availability of agile IT systems, capable of supporting this level of flexibility on the production network layer, as well as on the factory and process levels.

The FP7 project, ManuCloud, has been set up with the mission to investigate the production-IT related aspects of this transition and to develop and evaluate a suitable IT infrastructure to provide better support for on-demand manufacturing scenarios, taking multiple tiers of the value chain into account. On this path, ManuCloud seeks to implement the vision of a cloud-like architecture concept (see Figure 1). It provides users with the ability to utilize the manufacturing capabilities of configurable, virtualized production networks, based on cloud-enabled, federated factories, supported by a set of software-as-a-service applications.

Three industries have been selected to be the initial application context for the ManuCloud concepts and technologies: The photovoltaic (PV) industry, the organic lighting (organic light emitting diodes (OLED)) industry and the automotive supplies industry. Each industry is driven by specific market needs.

Over recent months, the market situation for the European PV industry has changed to a highly competitive environment. Prices for standard PV products have substantially dropped. China has significantly increased its market share while European companies have lost their leading position. This project will implement the ManuCloud infrastructure for the PV industry to evaluate whether highly customizable PV systems, especially in the area of building integrated photovoltaic, allow for new business models for this industry.

The market for organic lighting is in an earlier stage than the PV market. However, market research predicts the development of a multibillion dollar market for these products within a few years. Due to the unique properties of large-area diffuse light generation with

adjustable colors, organic lighting is expected to generate numerous new applications, a substantial share of which will be customized solutions. The project will set up and evaluate the ManuCloud infrastructure for customized organic lighting solutions.

In addition to these rather strategic applications, this project is expected to have an immediate impact on the automotive supplies industry, mainly on the factory/process level components of the ManuCloud infrastructure. The ability to add new functionalities to software systems at factory level and to quickly adjust production systems to new requirements is increasingly important for these companies. With typical state-of-the-art architectures used in production, additional functionality often causes an exponential growth of system complexity. This growth of complexity significantly increases ramp-up time, risk level, and costs as well as maintenance efforts for long-term operations.

Based on ManuCloud’s mission, two major R&D focal points have been selected for the project: The ManuCloud intra-factory environment

and the ManuCloud inter-factory environment. The intra-factory environment is comprised of production-related IT systems within a single factory which lays the foundation to connect the factory into the inter-factory environment. The inter-factory environment serves as a market place for virtualized manufacturing services, and supports the dynamic interconnection of multiple factories for specific purposes. For the intra-factory environment, the project intends to make heavy use of cross-fertilization effects in the area of best practices, standards and technologies available in the different industries represented by the project partners. The project will consider, among others, the standards families OPC-UA, SEMI (automation) and IEC61499.

The Unified Architecture (UA) is the next generation of the Open Connectivity standard that provides a cohesive, secure and reliable cross platform framework for access to real time and historical data and events. SEMI (Semiconductor Equipment and Materials International) is the global industry association serving the manufacturing supply chains for the micro-electronic, display and photovoltaic industries responsible for the generation

of standards specific to this area. IEC 61499 is a new standard of the International Electrotechnical Commission. It is event driven, enables engineering of complete, distributed systems and extensively supports hardware-independent engineering.

Special attention will be paid to the service interface of automation systems to the factory, including aspects of process capability modeling and system self description. A layer above the automation systems will support service discovery, management and orchestration, allowing for quick development and deployment of new factory-level services. The implementation of automation system services will be integrated with the engineering process for these systems.

The inter-factory environment supports a tightly controlled, on-demand integration of federated production-IT systems of different vendors, supporting joint specification management, shop-floor data transfer, high level of traceability and distributed quality management. This functionality will be provided by the ManuCloud Manufacturing as a Service (MaaS) environment. A front-end system will support the dynamic

configuration of virtual production networks and provide interfaces for product configurators, which are supported by a product design & manufacturing advisory subsystem. Demonstration scenarios will be setup for PV and OLED lighting use cases.

ManuCloud involves experts from eight different organizations that are directly included into the consortium and two additional third-party organizations from four different EU member states (Austria, Germany, Hungary and United Kingdom). From the direct consortium members, three organizations are SMEs, two are industry, two are research organizations and one is a university. The project will end in 2013.

**Links:**

ManuCloud project: <http://www.manucloud-project.eu>  
 Fraunhofer IPA: <http://www.ipa.fraunhofer.de>

**Please contact:**

Matthias Meier  
 Fraunhofer IPA, Germany  
 Tel: +49 711 970 1215  
 E-mail: [Matthias.Meier@ipa.fraunhofer.de](mailto:Matthias.Meier@ipa.fraunhofer.de)

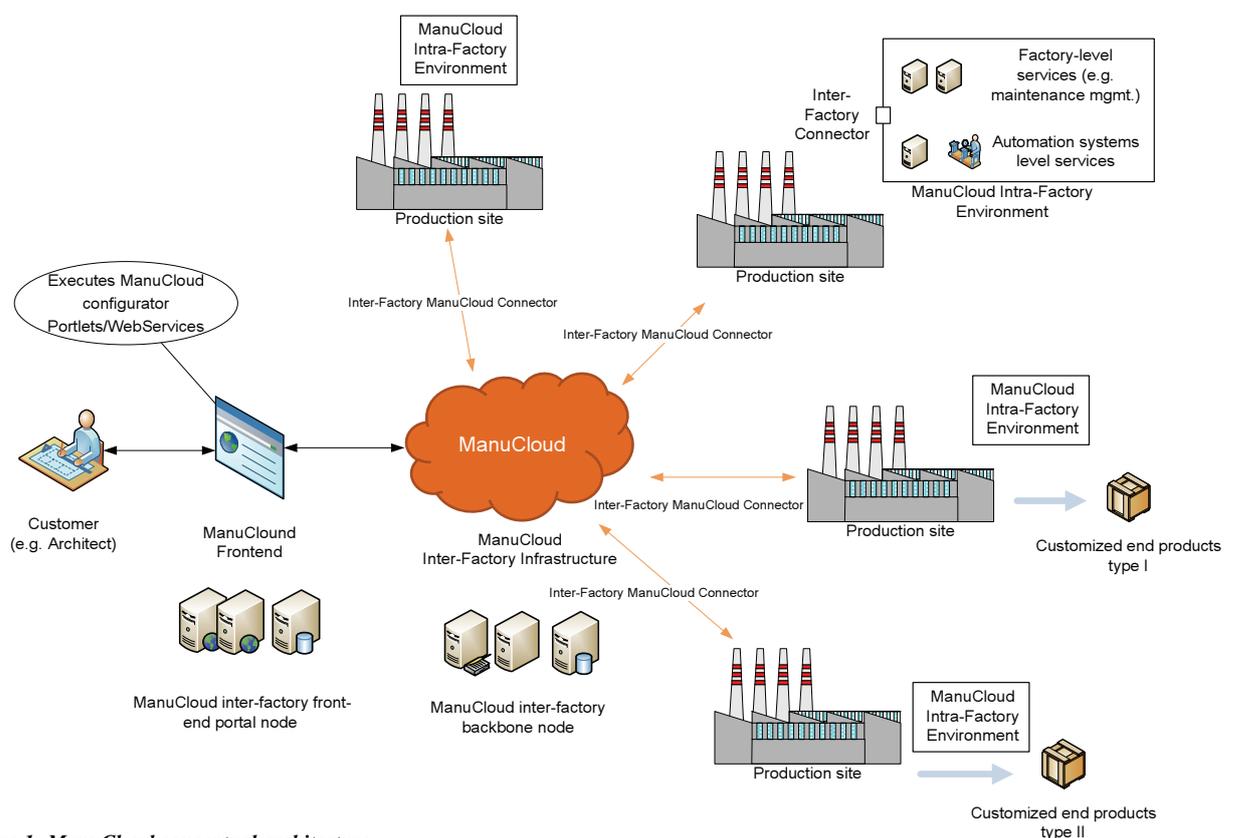


Figure 1: ManuCloud conceptual architecture.

# RESERVOIR – A European Cloud Computing Project

by Syed Naqvi and Philippe Massonet

The RESERVOIR project is developing breakthrough system and service technologies that will serve as an infrastructure as a service (IaaS) using Cloud computing. The project is taking virtualization forward to the next level in order to allow efficient migration of resources across geographies and administrative domains, maximizing resource exploitation, and minimizing their utilization costs.

Resources and Services Virtualization without Barriers (RESERVOIR) is a European Framework Programme 7 (FP7) funded project that will enable massive scale deployment and management of complex IT services across different administrative domains, IT platforms and geographies. Figure 1 shows a high-level description of RESERVOIR architecture. A detailed description is available on the project website. The project consortium is coordinated by IBM Haifa Research Lab, and includes a good balance of industry and academia.

The RESERVOIR project aims to support the emergence of Service-Oriented Computing (SOC) as a new computing paradigm. In this paradigm, services are software components exposed through network-accessible, platform and language independent interfaces, which enable the composition of complex distributed applications out of loosely coupled components. RESERVOIR project is extending, combining and integrating three core technologies: Virtualization, Grid computing and Business Service Management (BSM). This approach is taken in order to realise the vision of ubiquitous utility computing by harnessing the synergy between the complementary strengths of these technologies.

## Provisioning Services as Utilities

The vision of RESERVOIR is to enable the delivery of services on an on-demand basis, at competitive costs, and without requiring a large capital investment in infrastructure. This vision is inspired by the delivery of utilities in the physical world. A typical scenario in the physical world is the ability of an electrical grid in one country to dynamically provide more electric power to a grid in a neighbouring country to meet a spike in demand. It is evident that provisioning on-demand services from disparate service domains is far more complex than the provisioning of a utility in physical world. We are addressing these

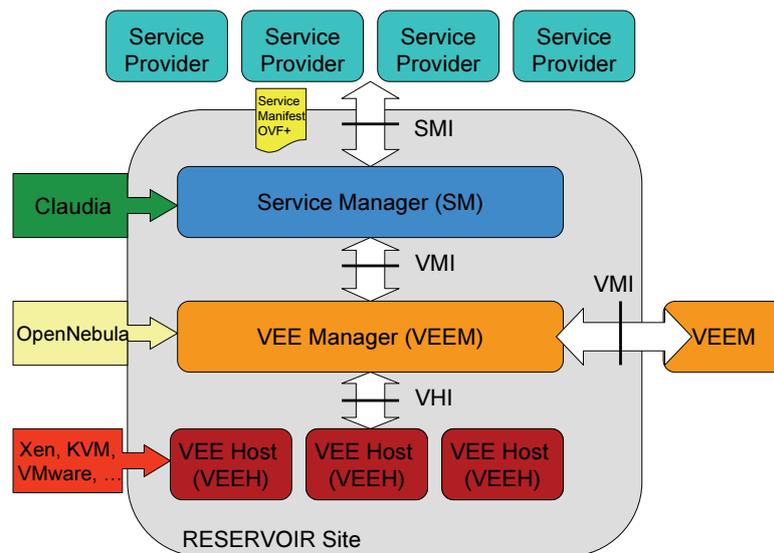


Figure 1: RESERVOIR architecture.

issues to overcome the barriers to delivering services as utilities.

## Advanced End-to-End Support for Service-Oriented Computing

Service-Oriented Computing (SOC) is a paradigm shift in the way software applications are designed and implemented to support business processes and users. SOC is built on the Service Oriented Architecture (SOA) that has the following characteristics:

- Services are reusable
- Services are autonomous, loosely coupled, and platform independent
- Services need to conform to the requirements of a service contract
- Services can be combined to form larger, multi-tier solutions
- Services can be described, published and discovered.

While SOA provides the general architecture to enable SOC, it does not address the fundamental issues required for an actual deployable solution, such as end-to-end security, service deployment, management and orchestration, service billing, and interpretation and

monitoring of Service Level Agreement (SLA) conditions. RESERVOIR, as an infrastructure project, is developing the technologies required to address these gaps, making the SOC paradigm a reality in the European economy.

## Service and Resource Migration without boundaries

The logical separation of a computing process and its physical hosting environment enables the migration of the process from one physical environment to another without affecting the process itself. Today, this separation is achieved either (1) by limiting the knowledge that the computing process has about its execution environment, or (2) by imposing hard configuration constraints to the infrastructure.

The first approach, commonly used in today's scientific applications, is restricted to self-contained independent processes that can be easily checkpointed and restarted elsewhere, for example processes that analyse large amounts of data but rarely interact with other processes or users. This restriction

renders this approach inappropriate for the type of commercial services that RESERVOIR aims to support.

In the second approach, widely used in today's commercial server virtualization offerings, the configuration of all the physical resources in the infrastructure has to be identical. For example, the commercial VMotion product from VMware only supports migration when the source and destination hypervisors are on the same subnet and have shared storage. Clearly, these configuration limitations make this approach inapplicable to large geographically distributed infrastructures that span several administrative domains.

RESERVOIR is removing these limitations/boundaries by taking virtualization forward to the next level enabling the migration of resources across geog-

raphies and administrative domains, maximizing resource exploitation, and minimizing their utilization costs.

#### Federated Heterogeneous Infrastructure and Management

Commercial virtualization systems typically offer non-standard management interfaces that are limited to their proprietary technologies and environments. RESERVOIR, in contrast, is developing an abstraction layer to facilitate the development of a set of high level management components that are not tied to any specific environment. To demonstrate the applicability of this generic management layer, we are using two different virtualization technologies: Virtual Machines (VMs) and Virtual Java Service Containers (VJSCs) by Sun. RESERVOIR is collaborating with standardization bodies (DMTF, LIBVIRT) to create standard interfaces that will enable inter-

action between distributed sites or Grid environments, allowing the federation of infrastructures.

#### Links:

<http://reservoir-fp7.eu>  
 Guide to RESERVOIR Framework:  
[http://www.reservoir-fp7.eu/uploads/Training/RESERVOIR\\_Framework\\_V1\\_022010.pdf](http://www.reservoir-fp7.eu/uploads/Training/RESERVOIR_Framework_V1_022010.pdf)  
[http://www.reservoir-fp7.eu/uploads/Training/RESERVOIR\\_Framework\\_Guide\\_Website.pdf](http://www.reservoir-fp7.eu/uploads/Training/RESERVOIR_Framework_Guide_Website.pdf)  
 CETIC: <http://www.cetic.be>

#### Please contact:

Syed Naqvi, CETIC, Belgium  
 Tel: +32 71 49 07 41  
 E-mail: [syed.naqvi@cetic.be](mailto:syed.naqvi@cetic.be)

Philippe Massonet, CETIC, Belgium  
 Tel: +32 71 49 07 44  
 E-mail: [philippe.massonet@cetic.be](mailto:philippe.massonet@cetic.be)

## Managing Virtual Resources: Fly through the Sky

by Jérôme Gallard and Adrien Lèbre

*Virtualization technologies have been a key element in the adoption of Infrastructure-as-a-Service (IaaS) cloud computing platforms as they radically changed the way in which distributed architectures are exploited. However, a closer look suggests that the way of managing virtual and physical resources still remains relatively static.*

Through an encapsulation of software layers into a new abstraction - the virtual machine (VM) - cloud computing users can run their own execution environment without considering, in most cases, software and hardware restrictions as it was formerly imposed by computing centers. By using specific APIs, users can create, configure and upload their VM to IaaS providers, which in turn are in charge of deploying and running the requested VMs on their physical architecture.

Due to the growing popularity of these IaaS platforms, the number of VMs and consequently the amount of data to manage is increasing. Consequently, IaaS providers have to permanently invest in new physical resources to extend their infrastructure. This leads to concerns, which are being addressed by the Cluster and Grid Communities, about the management of large-scale distributed resources.

Similarly to previous works that led to the Grid paradigm, several works propose the interconnection of distinct IaaS platforms to produce a larger infrastructure. In cloud computing terminology, we refer to such a federation as a sky computing platform. Although it improves flexibility in terms of VM management by delivering, for instance, additional resources to users when one site is overloaded or by offering more efficient nodes when relevant, most of the available IaaS solutions assign the VMs to the physical machines in a static manner and without reconsidering the allocation throughout the whole infrastructure during their execution.

Such an allocation strategy is not appropriated to tackle volatility constraints intrinsic to large-scale architectures (node additions/removals, node/network failures, energy constraints etc.). Furthermore, it does not allow fine-

scale management of resource assignment according to the VM's fluctuating needs. These two concerns have been the focus of the 'Saline' and 'Entropy' projects, three year projects that aim to manage virtualization environments more dynamically across distributed architectures. Both rely on an encapsulation of each computing task into one or several VMs according to the nature of the task.

Carried out by the MYRIADS team from the INRIA research center in Rennes, France, the Saline proposal focuses on aspects of grid volatility. Through periodic snapshots and a monitoring of the VMs of each task, Saline is able to restart the set of VMs that may malfunction due to a physical failure by resuming latest VM snapshots. Keeping in mind that Saline is still dealing with some restrictions (such as external communications) and considering that a Saline manager is deployed on each

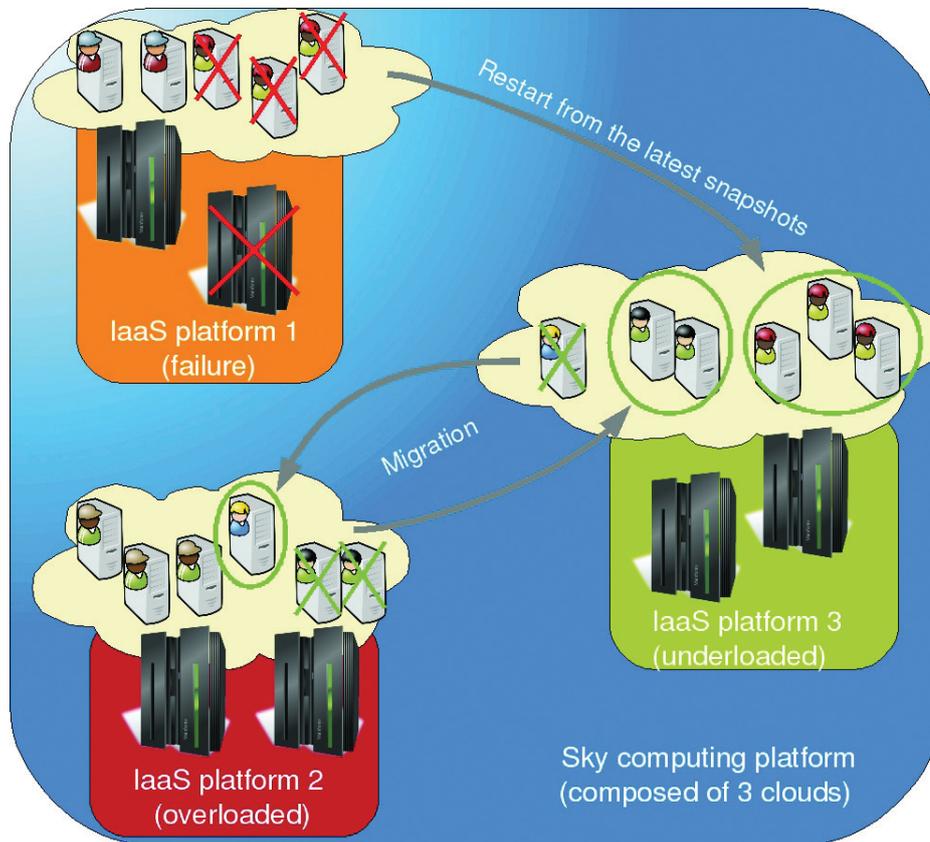


Figure 1: Sky computing platforms composed of three clouds.

site, the resume operations can be performed anywhere in the grid thanks to advanced management of the network configuration of VMs.

Developed by the ASCOLA research group, a joint team between INRIA and the Ecole des Mines de Nantes, France, Entropy is a virtual machine manager for clusters. It acts as an infinite control loop, which performs cluster-wide context switches (i.e. permutation between active and inactive VMs present in one cluster) to provide a globally optimized placement of VMs according to the real usage of physical resources and the scheduler objectives (consolidation, load-balancing etc.). The major advantage concerns the cluster-wide context switch operation that is performed in a minimum number of actions and in the most efficient way.

The integration of both projects to develop a unique solution started during summer 2010. Our main objective is to combine available mechanisms provided by both systems to deliver an advanced management of virtualized environments across a sky computing platform. Once the VMs are created and uploaded somewhere into the sky, cloud computing users will let this new framework manage their environment

throughout the different sites: each set of VMs may “fly” from one cloud to another one according to the allocation policy and the physical changes of the sky. Technically speaking, Saline focuses on the transfer and the reconfiguration of the VMs across the sky, whereas Entropy is in charge of efficiently managing VMs on each cloud. A first prototype is under development and preliminary experiments have been carried out on the Grid’5000 testbed.

Future work will complete the implementation with additional mechanisms provided by virtualization technologies such as emulation and aggregation in order to improve relocation possibilities in the sky, which are currently limited by physical considerations such as processor architecture, size of memory, etc. Finally, our long-term objective is to completely dissociate the vision of resources that each cloud computing user expects from the physical one delivered by the sky. Such a framework will deliver the interface between these two visions by providing first a description language capable of representing user expectations in terms of VMs and second, adequate mechanisms enabling setup and maintenance of this virtual vision in case of physical changes as previously discussed.

#### Links:

ASCOLA research group:  
<http://www.emn.fr/x-info/ascola>  
 INRIA MYRIADS team:  
<http://www.irisa.fr/myriads>  
 Entropy project:  
<http://entropy.gforge.inria.fr>

#### Please contact:

Jérôme Gallard  
 INRIA, France  
 Tel: +33 299 842 556  
 E-mail: [jerome.gallard@inria.fr](mailto:jerome.gallard@inria.fr)

Adrien Lèbre  
 Ecole des Mines de Nantes, France  
 Tel: +33 251 858 243  
 E-mail: [adrien.lebre@mines-nantes.fr](mailto:adrien.lebre@mines-nantes.fr)

# OW2 ProActive Parallel Suite: Building Flexible Enterprise CLOUDs

by Denis Caromel, Cédric Dalmaso, Christian Delbe, Fabrice Fontenoy and Oleg Smirnov

*The ProActive Parallel Suite features Infrastructure as a Service (IaaS) capabilities together with an innovative parallel programming model and a distributed workflow environment. It involves the OASIS team from INRIA Sophia Antipolis, which initiated the development early 2000 and ActiveEon, an INRIA spin-off created in 2007, which together co-develop ProActive and provide users with professional services.*

Federating large sets of distributed resources is an important issue for companies. Scientists and engineers in many fields including finance, engineering, entertainment, and energy need increasing amounts of computational power. As a consequence, companies and laboratories are placing increasing demands on their existing infrastructure. Due to peak workloads, companies also want to gain flexibility. Finally, green IT with improved control of infrastructure usage is another reason for desiring more precise control with optimization of the overall workload. In order to address these strong industrial and scientific needs INRIA and ActiveEon provide IT manager with a simple way to aggregate native or virtualized resources.

## Resource Provisioning

ProActive Resourcing is the first building block used to provide heterogeneous resource management. ProActive Resourcing provides an open source intelligent and adaptive application deployment engine to virtualize hardware resources and monitor and control all computing resources. With ProActive, resource management is easier and highly configurable. It leverages the existing infrastructure of an organization, from dedicated clusters to heterogeneous distributed resources, building a Private Cloud with the capacity to manage virtualization and software appliances.

We introduce the concept of Node Source which associates the method of acquisition of computing resources with the policy determining when these resources have to deploy or to be released. Node sources allow a company to have business driven management of their computing power. New resources can be automatically acquired at any given time, and acquisition can be automatically triggered according to the load

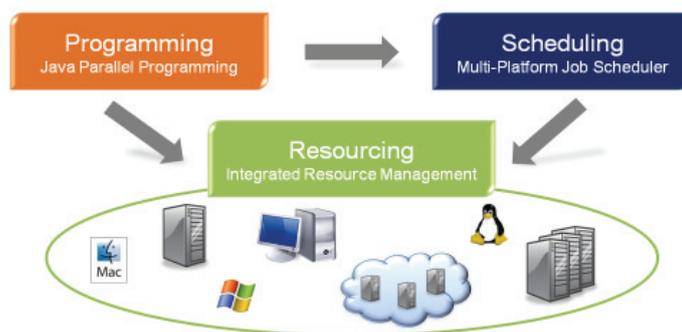


Figure 1: ProActive Parallel Suite.

of the current infrastructure. These new resources can come from another business unit in the same company or from outside the company like a Data Center or a public cloud.

The solution, developed in Java, is highly portable and can be deployed on Unix-like, Windows, and Mac operating systems. Resources can be virtualized or not. The following virtualization environments are supported: VMware, KVM, Xen, Xen Server, QMU, and Microsoft Hyper-V. In case of native resources directly accessed through well known protocols, RSH and SSH can be used, as well as third party schedulers like PBS, LSF, SGE, IBM Load Leveler, Oar, Prun.

## Job Scheduling and Workload Management

ProActive Scheduling is an open source multi-platform job scheduler managing the distribution of workflows and application execution over the available computing resources. ProActive Scheduling ensures more work is done with fewer resources: maximum utilization and optimal allocation of existing IT infrastructure, reducing administration costs and future hardware expenditures.

A command line interface (CLI) as well as a graphical user interface (GUI)

based on Eclipse RCP provides users and administrators with all the tools needed to easily submit, monitor and retrieve results, and administrate enterprise cloud. Moreover, workflows can be built using various methods such as XML file, flat file, Web Service and programming API in Java and C/C++. The ability to use a range of languages facilitates integration with any kind of application when part of a workload needs to be delegated to other resources.

## Operating a production Platform: ProActive PACA Grid

The ProActive PACA Grid is a Computing Cloud operated by INRIA, University of Nice and CNRS-I3S laboratory; it is funded by the PACA Lander and the European Commission. The Cloud platform makes available a set of machines to laboratories and SMEs. The resources are accessible via graphical interactive interfaces launchable from the PACA Grid website, in a portal mode.

The machines are currently deployed within INRIA Sophia Antipolis networks. The Cloud aggregates dedicated machines, both Linux and Windows, GPU processors, and spare desktop machines, dynamically added during nights and week-ends. Infrastructure and workload are managed using ProActive Resourcing Scheduling. It is

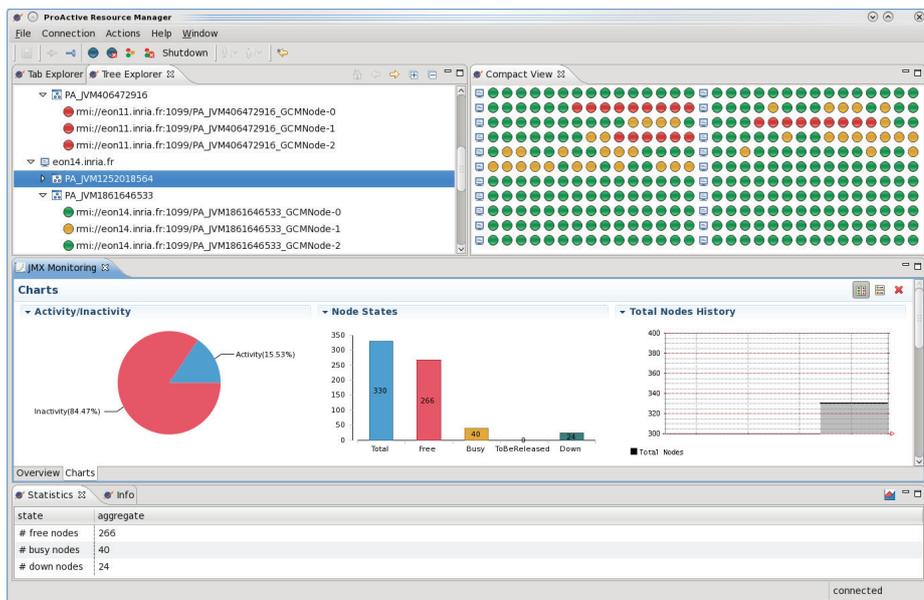


Figure 2: ProActive Resource Manager.

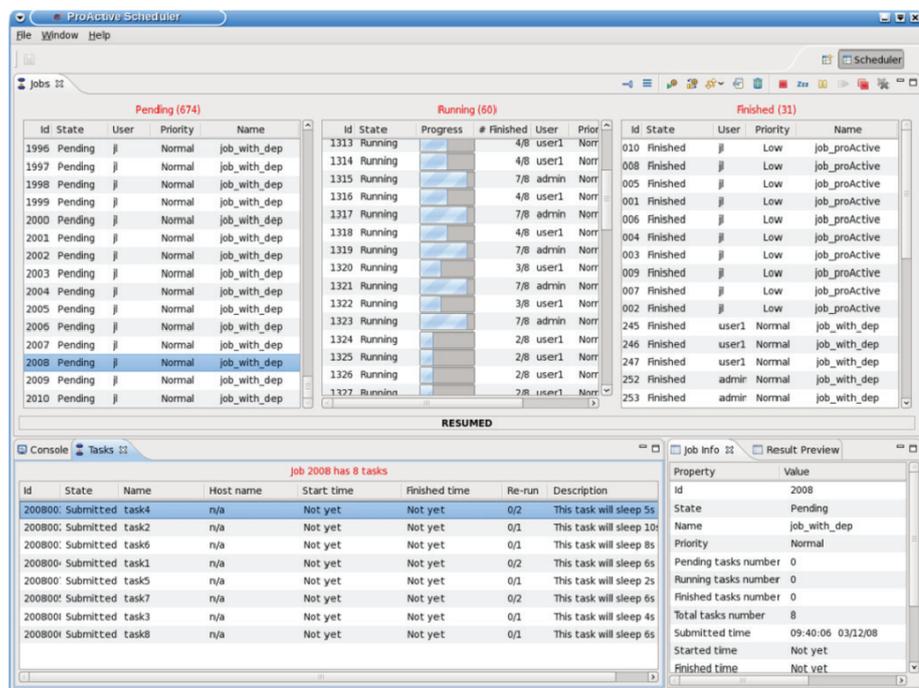


Figure 3: ProActive Scheduler.

integrated with the infrastructure through JMX/Nagios for the monitoring and LDAP for the authentication. The use of PACA Grid is simplified by the DataSpaces feature which automatically transfers input file parameters and brings home output results. Today, ProActive PACA Grid features in production over 1000 CPU Cores, 4 TByte of RAM, 15 TByte of storage, and 480 CUDA Cores for about 2 TFlops.

#### Conclusion and Perspectives

The comprehensive ProActive Parallel Suite toolkit will be soon enriched with a graphical editor for Workflows allowing design and monitoring of jobs of tasks, and a Web Portal to benefit

from thin client interfaces. In addition, ProActive is one of the three building blocks of the Open Source OW2 Cloud Initiative (OSCI) recently launched by the consortium.

The OASIS has collaborated with many EU partners such as University of Pisa, IBM, Atos Origin, Thales, Microsoft, HP, Oracle and Telefónica, including co-operation with other ERCIM members including CNR and Fraunhofer, and EU projects GridCOMP, CoreGRID, SOA4ALL, and TEFIS. The ProActive team has also built close relationships with international partners like University of Adelaide, Tsinghua University, University of Chile, STIC in Shanghai.

#### Links:

ActiveEon SAS:

<http://www.activeeon.com/>

ProActive Parallel Suite

<http://proactive.inria.fr/>

PACA Grid:

<http://proactive.inria.fr/pacagrid/>

OW2 OSCI:

<http://www.ow2.org/view/Cloud/>

#### Please contact:

Denis Caromel,

INRIA Sophia Antipolis, France

Tel: +33 4 92 38 76 31

E-mail: [Denis.Caromel@inria.fr](mailto:Denis.Caromel@inria.fr)

# FoSII - Foundations of Self-Governing ICT Infrastructures

by Vincent C. Emeakaroha, Michael Maurer, Ivona Brandic and Schahram Dustdar

*The DSG Group at Vienna University of Technology is investigating self-governing Cloud Computing infrastructures necessary for the attainment of established Service Level Agreements (SLAs). Timely prevention of SLA violations requires advanced resource monitoring and knowledge management. In particular, we develop novel techniques for mapping low-level resource metrics to high-level SLAs, monitoring resources at execution time, and applying Case Based Reasoning for the prevention of SLA violations before they occur while reducing energy consumption, ie, increasing energy efficiency.*

Cloud computing is a promising technology for the realization of large, scalable on-demand computing infrastructures. Currently, many enterprises are adopting this technology to achieve high performance and scalability for their applications while maintaining low cost. Service provisioning in the Cloud is based on a set of predefined non-functional properties specified and negotiated by means of Service Level Agreements (SLAs). Cloud workloads are dynamic and change constantly. Thus, in order to reduce steady human interactions, self-manageable Cloud techniques are required to comply with the agreed customers' SLAs.

Flexible and reliable management of SLAs is of paramount importance for both Cloud providers and consumers. On the one hand, the prevention of SLA violations avoids penalties that are costly to providers. On the other hand, based on flexible and timely reactions to possible SLA violation threats, user interaction with the system can be minimized

enabling Cloud computing to take roots as a flexible and reliable form of on-demand computing. Furthermore, a trade-off has to be found between proactive actions that prevent SLA violations and those that reduce energy consumption, ie, increase energy efficiency.

The Foundation of Self-governing ICT Infrastructures (FoSII) research project is proposing solutions for autonomic management of SLAs in the Cloud. The project started in April 2009 and is funded by the Vienna Science and Technology Fund (WWTF). In this project, we are developing models and concepts for achieving adaptive service provisioning and SLA management via resource monitoring and knowledge management techniques.

Figure 1 depicts the components of the FoSII infrastructure. Each FoSII service implements three interfaces: (i) negotiation interface necessary for the establishment of SLA agreements, (ii) service management interface neces-

sary for starting service, uploading data, and similar management actions, and (iii) self-management interface necessary to devise actions in order to prevent SLA violations.

The self-management interface as shown in Figure 1 specifies operations for sensing changes of the desired state and for reacting to those changes. The host monitor sensors continuously monitor the infrastructure resource metrics (input sensor values arrow a in Figure 1) and provide the knowledge component with the current resource status. The run-time monitor sensors sense future SLA violation threats (input sensor values arrow b in Figure 1) based on resource usage experiences and predefined thresholds.

As shown in Figure 1, the Low-level Metric to High-level SLA (LoM2His) framework is responsible for monitoring and sensing future SLA violation threats. It comprises the host monitor and the run-time monitor. The host monitor monitors low-level resource metrics such as CPU, memory, disk space, incoming bytes, etc using monitoring agents like Gmond from Ganglia project embedded in each Cloud resource. It extracts the monitored output from the agents, processes them and sends the metric-value pairs through our implemented communication model to the run-time component.

The run-time component receives the metric-value pairs and, based on predefined mapping rules, maps them into equivalent high-level SLA parameters. An example of an SLA parameter is service availability  $A_v$ , which is calculated using the resource metrics downtime and uptime as follows:  

$$A_v = (1 - \text{downtime}/\text{uptime}) \times 100.$$
 The provider defines the mapping rules using appropriate Domain Specific

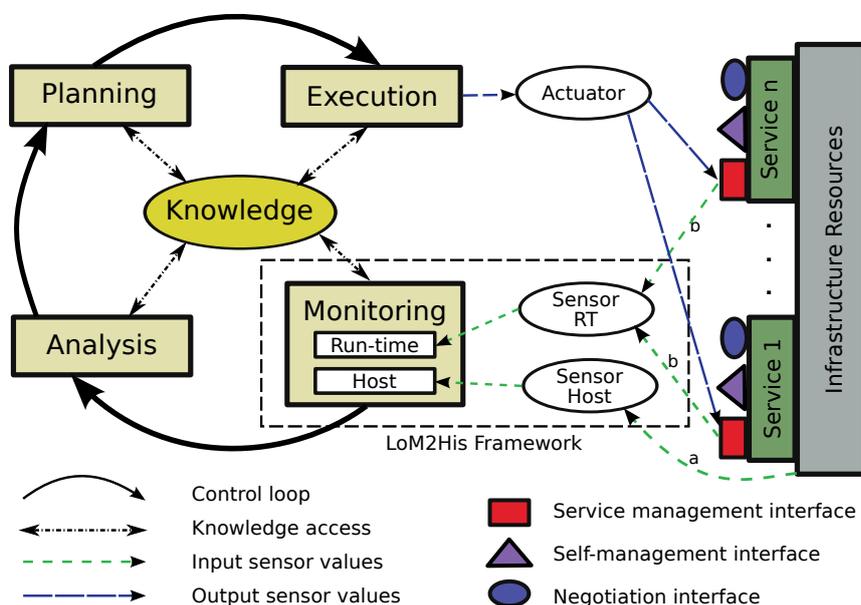


Figure 1: FoSII Infrastructure.

Languages (DSLs). The concept of detecting future SLA violation threats is designed by defining a more restrictive threshold than the SLA violation thresholds known as threat threshold. Thus, calculated SLA values are compared with the predefined threat threshold in order to react before an SLA violation occurs. In case SLA violation threats are detected, the run-time monitor sends notification messages to the knowledge component for preventive actions.

During the analysis and planning phases the knowledge component then suggests appropriate actions to solve SLA violation threats. As a conflicting goal, it also tries to reduce energy consumption by removing resources from over-provisioned services. Reactive actions thus include increasing or decreasing memory, storage or CPU usage for each service. After the action

has been executed the knowledge component learns the utility of the action in this specific situation via Case Based Reasoning (CBR). CBR contains previously solved cases together with their actions and utilities, and tries to find the most similar case with the highest utility for each new case. Furthermore, it examines the timing and the effectiveness of an action, ie, whether the action would have helped but was triggered too late, or was unnecessarily triggered too early, and consequently, it updates the threat thresholds from the monitoring component. In the future, the knowledge component will offer different energy efficiency classes that will reflect the trade-off between preventing violations and saving energy, and it will integrate knowledge about penalties and client's status for prioritizing resource demand requests when resources are scarce.

We have successfully implemented the first versions of the LoM2HiS framework and the knowledge component. First evaluation results of the components have been published in top-ranked international conferences: HPCS 2010, COMPSAC 2010, SERVICES 2010, and CloudComp 2010.

**Links:**

<http://www.infosys.tuwien.ac.at/linksites/FOSII/index.html>

<http://www.infosys.tuwien.ac.at/>

<http://www.infosys.tuwien.ac.at/staff/vincent/>

**Please contact:**

Vincent Chimaobi Emeakaroha

Vienna University of Technology /

AARIT, Austria

Tel +43 1 58801 18457

E-mail: [vincent@infosys.tuwien.ac.at](mailto:vincent@infosys.tuwien.ac.at)

## Large-Scale Cloud Computing Research: Sky Computing on FutureGrid and Grid'5000

by Pierre Riteau, Maurício Tsugawa, Andréa Matsunaga, José Fortes and Kate Keahey

*How can researchers study large-scale cloud platforms and develop innovative software that takes advantage of these infrastructures? Using two experimental testbeds, FutureGrid in the United States and Grid'5000 in France, we study Sky Computing, or the federation of multiple clouds.*

The remarkable development of cloud computing in the past few years, and its proven ability to handle web hosting workloads, is prompting researchers to investigate whether clouds are suitable to run large-scale scientific computations. However, performing these studies using available clouds poses significant problems. First, the physical resources are shared with other users, which can interfere with performance evaluations and render experiment repeatability difficult. Second, any research involving modification of the virtualization infrastructure (eg, hypervisor, host operating system, or virtual image repository) is impossible. Finally, conducting experiments with a large number of resources provisioned from a commercial cloud provider incurs high financial cost, and is not always possible due to limits to the maximum number of resources one can use. These problems, which would have been limitations for our sky computing experiments, were avoided by our use of experimental testbeds.

We study sky computing, an emerging computing model where resources from multiple cloud providers are leveraged to create large-scale distributed virtual clusters. These clusters provide resources to execute scientific computations requiring large computational power. Establishing a sky computing system is challenging due to differences among providers in terms of hardware, resource management, and connectivity. Furthermore, scalability, balanced distribution of computation, and measures to recover from faults are essential for applications to achieve good performance.

Experimental distributed testbeds offer an excellent infrastructure to carry out our sky computing experiments. We make use of the following two platforms: FutureGrid, a new experimental grid testbed distributed over six sites in the United States, and Grid'5000, an infrastructure for large-scale parallel and distributed computing research

composed of nine sites in France. Using the reconfiguration mechanisms provided by these testbeds, we are able to deploy the Nimbus open source cloud toolkit on hundreds of nodes in a matter of minutes. This gives us exclusive access to cloud platforms similar to real-world infrastructures, such as Amazon EC2. Full control of the physical resources and of their software stack guarantees experiment repeatability. Combining two testbeds gives us access to more resources and, more importantly, offers a larger geographical distribution, with network latencies and bandwidth on a par with those found on the Internet. Our project is the first combining these two testbeds, paving the way for further collaboration.

Several open source technologies are integrated to create our sky computing infrastructures. Xen (an open source platform for virtualization) machine virtualization is used to minimize platform (hardware and operating system

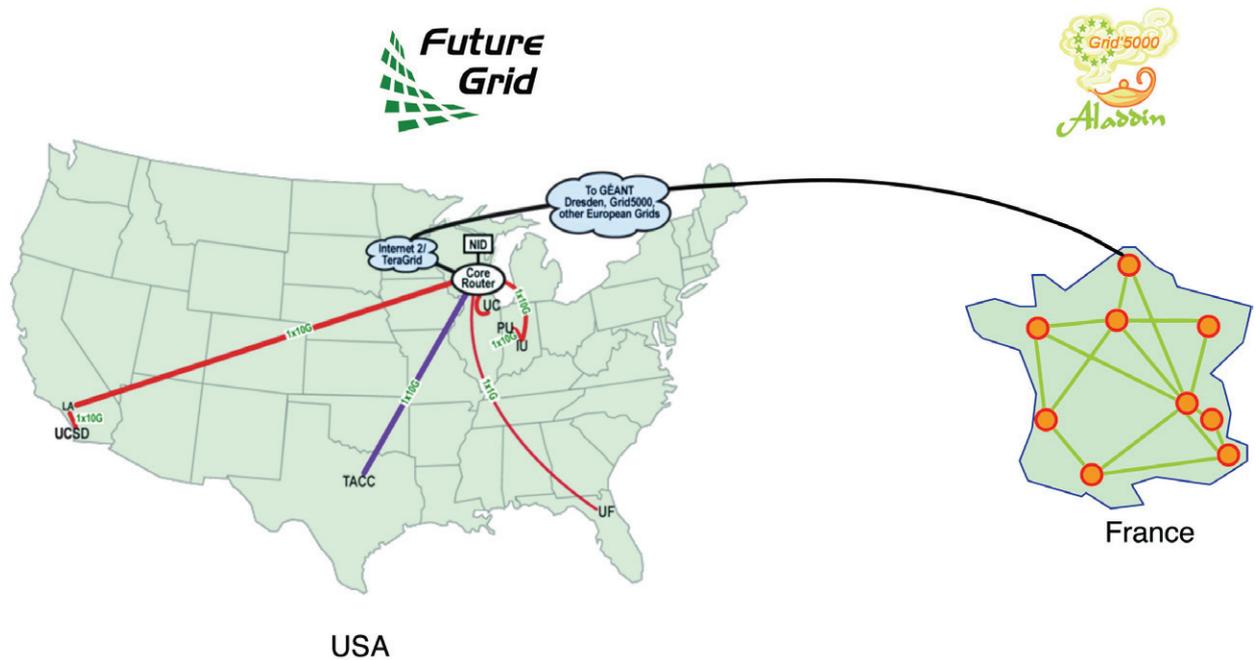


Figure 1: The FutureGrid and Grid'5000 testbeds used for Sky Computing research.

stack) differences. Nimbus, which provides both an Infrastructure-as-a-Service implementation with EC2/S3 interfaces and higher-level cloud services such as contextualization, is used for resource and virtual machine (VM) management. By deploying Nimbus on FutureGrid and Grid'5000, we provide an identical, Amazon Web Services (AWS)-compatible, interface for requesting virtual resources on these different testbeds, making interoperability possible. We then leverage the contextualization services offered by Nimbus to automatically configure the provisioned virtual machines into a virtual cluster without any manual intervention.

Commercial clouds and scientific testbeds limit the network connectivity of virtual machines making all-to-all communication, required by many scientific applications, impossible. 'ViNe', a virtual network based on an IP-overlay, allows us to enable all-to-all communication between virtual machines involved in a virtual cluster spread across multiple clouds. In the context of FutureGrid and Grid'5000, it allows us to connect the two testbeds with minimal intrusion in their security policies. Once the virtual cluster is provisioned, we configure it with Hadoop (open-source software for distributed computing) to provide a platform for parallel fault-tolerant execution of a popular embarrassingly parallel bioinformatics application (BLAST). We

further leverage the dynamic cluster extension feature of Hadoop to experiment with the addition of new resources to virtual clusters as they become available. New virtual resources from Grid'5000 and FutureGrid are able to join the virtual cluster while computation is under progress. As resources are added, map and reduce tasks are distributed among these resources, speeding up the computation process. To accelerate the provisioning of additional Hadoop worker virtual machines, we developed an extension to Nimbus taking advantage of Xen copy-on-write image capabilities. This extension decreases the VM instantiation time from several minutes to a few seconds.

Ongoing and future activities involve solving scalability issues in cloud computing infrastructures when requesting resources for large-scale computation, allowing transparent elasticity of sky computing environments, and using live migration technologies to take advantage of dynamicity of resources between multiple cloud platforms.

Our project is a collaboration started in 2010 between the Myriads research team at the IRISA/INRIA Rennes - Bretagne Atlantique laboratory in France, the ACIS laboratory at University of Florida, and the Computation Institute at the Argonne National Laboratory and the University of Chicago.

#### Links:

<http://futuregrid.org/>  
<https://www.grid5000.fr/>  
<http://www.nimbusproject.org/>

#### Please contact:

Pierre Riteau,  
 Université de Rennes 1, IRISA/INRIA  
 Rennes, France  
 Tel: +33 2 99 84 22 39  
 E-mail: [Pierre.Riteau@irisa.fr](mailto:Pierre.Riteau@irisa.fr)

# elasticLM – Software License Management for Distributed Computing Infrastructures

by Claudio Cacciari, Daniel Mallmann, Csilla Zsigri, Francesco D’Andria, Björn Hagemeier, Angela Rimpl, Wolfgang Ziegler and Josep Martrat

**One of the major obstacles to using commercial applications in Distributed Computing Infrastructures like Grids or Clouds is the current technology that relies on controlling the use of these applications with software licenses. “Software licensing practices are limiting the acceleration of grid adoption” was one of the results of a survey of the 451group in 2005. Just recently the 451group published a similar report on obstacles to the broad adoption of Cloud Computing - and again licensing practices were listed among the top five obstacles. elasticLM overcomes the limitations of existing licensing technologies allowing seamless running of license protected applications in computing environments ranging from local infrastructures to external Grids and Clouds.**

So far, commercial software has rarely been used in Grids and is – except for Software as a Service (SaaS) – also rarely used in “public” Clouds. This is due to current limitations of license management technology and the fact that business models of independent software vendors (ISV) do not allow for the use of their software in the Grid.

Only recently MathWorks provided a technical solution (and a business

required licenses. Usually, these licenses are provided on the basis of named users, IP-addresses, or sometimes as a site license. Executing these applications is almost impossible or illegal when using resources that are spread across different administrative domains. Licenses are usually bound to the license server within the domain of the user which does not allow access for verification from outside, due to firewalls or similar.

menting software licenses as Web Service resources, thus providing platform independent access to licenses just like to other virtualized resources, is at the core of the SmartLM architecture. Licenses are managed through a license service implemented as a bag of services and realised as mobile tokens, delivering the required flexibility and mobility. In order to execute a license protected application under current technology, a permanent bi-directional network connection to the license service at runtime controlling the authorization is required. elasticLM aims to solve this problem by decoupling authorization for license usage from authorization for application execution.. All authorizations for license usage are expressed and guaranteed by Service Level Agreements based on the WS-Agreement specification of the Open Grid Forum (OGF).

A built-in license scheduling service allows licenses to be reserved in advance, which means license availability can be guaranteed at a given later point in time, eg when the computing resources to execute the simulation become available. Thus there is no risk of encountering blocked licenses while an application is idling waiting for computing resources to become available. Similarly, there is no risk of aborted applications because the required license is used by another user at the time the application starts up after waiting for resources. An orchestration service can be used to synchronize license reservation and resource availability.

Using open standards as far as possible, instead of proprietary protocols, is considered crucial for the interoperability with and integration into

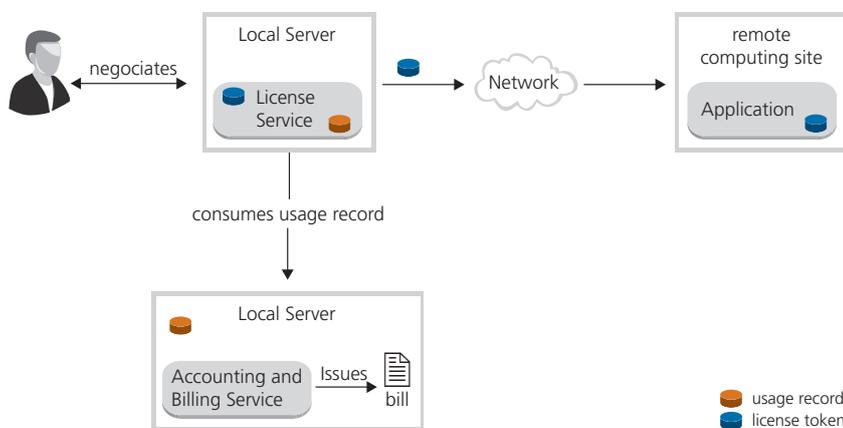


Figure 1: Basic scenario for using elasticLM licenses in a distributed computing infrastructure.

model) allowing use of their MATLAB suite in the EGEE (Enabling Grids for E-science project) Grid. However, this is a bilateral agreement only and has no implications for other Grids. IBM recently achieved an agreement on providing licenses for certain IBM applications inside EC2 with Amazon. Apart from these exceptions however, license management technology for software is still based on the model of local computing centres providing both resources for computation and the software used for simulations together with the

The European project SmartLM laid the foundations for elasticLM. After the SmartLM project ended in July 2010, three of the project’s partners (Atos Origin, Gridcore Aktiebolag and the Fraunhofer Institute SCAI) agreed to develop the new product based on the SmartLM prototype.. The SmartLM approach reflects the changing paradigms within information technology. The transition from monolithic infrastructures to agile, service-based architectures was at the heart of the SmartLM project. Treating and imple-

existing middleware stacks. Figure 1 shows the basic SmartLM scenario for environments where at run-time there is no bi-directional network link available between the license service that created the license token and the remote execution environment, eg due to firewall restrictions.

SmartLM addresses the licensing management issues not only from a technological point of view, but also from the perspective of developing new business models. This approach is necessary in order to convince Independent Software

Vendors adopting to the new license technology. More details can be found on the project web pages.

The major part of the licensing technology presented in this article has been designed and implemented prototypically in the European Commissions ICT programme in the FP7 project SmartLM. In the European funded project OPTIMIS (Optimized Infrastructure Services) we will further improve the capabilities of the SmartLM solution, developing additional features, for instance a feature

that makes the SmartLM solution more secure in Clouds or extending the capabilities if there is a bi-directional network connection available at run-time.

#### Links:

elasticLM <http://www.elasticlm.com>  
SmartLM <http://www.smartlm.eu/>

#### Please contact:

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

## Enabling Reliable MapReduce Applications in Dynamic Cloud Infrastructures

by Fabrizio Marozzo, Domenico Talia and Paolo Trunfio

*MapReduce is a parallel programming model for large-scale data processing that is widely used in Cloud computing environments. Current MapReduce implementations are based on master-slave architectures that do not cope well with dynamic Cloud infrastructures, in which nodes join and leave the network at high rates. We have designed a MapReduce architecture that uses a peer-to-peer approach to manage node churn and failures in a decentralized way, so as to provide a more reliable MapReduce middleware that can be effectively exploited in dynamic Cloud infrastructures.*

MapReduce is a framework for processing large data sets in a highly parallel way by exploiting computing facilities available in a data centre or through a Cloud computing infrastructure. Programmers define a MapReduce application in terms of a map function that processes a key/value pair to generate a list of intermediate key/value pairs, and a reduce function that merges all intermediate values associated with the same intermediate key.

Current MapReduce implementations, like Google's MapReduce, are based on a master-slave architecture. A job is submitted by a user node to a master node that selects idle workers and assigns a map or reduce task to each. When all the tasks have been completed, the master node returns the result to the user node. The failure of one worker is managed by re-executing its task on another worker, while master failures are not explicitly managed as designers consider failures unlikely in reliable computing environments, such as a data centre or a dedicated Cloud.

In contrast, node churn and failures – including master failures – are likely in

dynamic Cloud environments, such as a Cloud of clouds, which can be formed by a large number of computing nodes that join and leave the network at very high rates. Therefore, providing effective mechanisms to manage such problems is fundamental to enable reliable MapReduce applications in dynamic Cloud infrastructures, where current MapReduce middleware could be unreliable.

At the University of Calabria and ICAR-CNR we have designed an adaptive MapReduce framework, called P2P-MapReduce, which exploits a peer-to-peer model to manage node churn, master failures, and job recovery in a decentralized but effective way, so as to provide a more reliable MapReduce middleware that can be effectively exploited in dynamic Cloud infrastructures.

P2P-MapReduce exploits the peer-to-peer paradigm by defining an architecture in which each node can act either as a master or a slave. The role assigned to a given node depends on the current characteristics of that node, and can change dynamically over time. Thus, at

each time, a limited set of nodes is assigned the master role, while the others are assigned the slave role. Each master node acts as a backup node for the other master nodes. A user node can submit a job to one of the master nodes, which will manage it as usual in MapReduce. That master dynamically replicates the entire job state (ie, the assignments of tasks to nodes, the locations of intermediate results, etc.) on its backup nodes. If those backup nodes detect the failure of the master, they will elect one of them as a new master that will manage the job computation using its local replica of the job state.

The behaviour of a generic node is modelled as a state diagram which defines the different states that a node can assume, and all the events that determine transitions from one state to another state (see Figure 1). The slave macro-state describes the behaviour of an active or idle worker. The master macro-state is modelled with three parallel states, which represent the different roles a master can perform concurrently: possibly acting as a primary master for one or more jobs (management); possibly acting as a backup

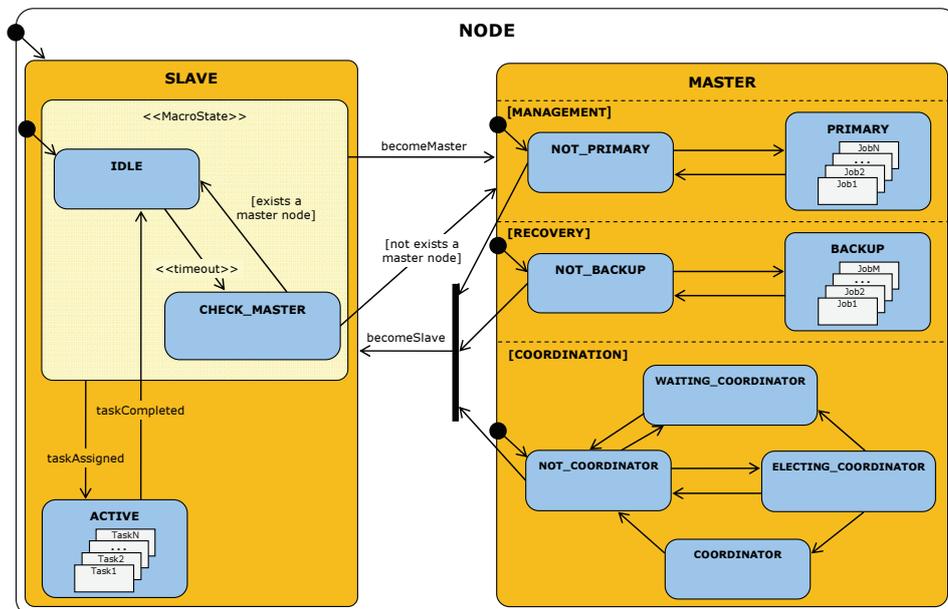


Figure 1: UML State Diagram describing the behavior of a generic node in the P2P-MapReduce framework. The slave macro-state describes the behavior of an active or idle worker. The master macro-state is modelled with three parallel states: Management (the node is possibly acting as a primary master); Recovery (the node is possibly acting as a backup master); Coordination (the node is possibly acting as the network coordinator).

master for one or more jobs (recovery); coordinating the network (coordination). The goal of a master acting as the network coordinator is to ensure the presence of a given percentage of masters on the total number of nodes; to this end, it has the power to change slaves into masters, and vice versa.

We implemented a prototype of the P2P-MapReduce framework using the Sun's JXTA peer-to-peer framework. In our implementation, each node includes three software modules/layers: Network, Node and MapReduce (see Figure 2). The Network module is in charge of the interactions with the other nodes using the pipe communication mechanism provided by the JXTA framework; additionally, it allows the node to interact with the JXTA Discovery Service for publishing its features and for querying the system (eg, when looking for idle slave nodes). The Node module controls the node lifecycle; its core is represented by the FSM component which implements the logic of the finite state machine shown in Figure 1. Finally, the MapReduce module manages the local execution of jobs (when the node is acting as a master) or tasks (when the node is acting as a slave). Currently this module is built upon the local execution engine of Apache Hadoop.

We are carrying out a set of experiments to evaluate the behaviour of the P2P-MapReduce framework compared to a standard master-slave implementation of MapReduce, in the presence of different levels of churn. Early experimental results show that, in contrast to

standard implementations, the P2P-MapReduce framework does not suffer from job failures even in presence of very high churn rates, thus enabling the execution of reliable MapReduce applications in very dynamic Cloud infrastructures.

#### Links:

<http://labs.google.com/papers/mapreduce.html>  
<https://jxta.dev.java.net>  
<http://hadoop.apache.org>

#### Please contact:

Domenico Talia  
 ICAR-CNR and  
 DEIS, University of Calabria, Italy  
 Tel: +39 0984 494726  
 E-mail: [talia@deis.unical.it](mailto:talia@deis.unical.it)

Fabrizio Marozzo and Paolo Trunfio  
 DEIS, University of Calabria, Italy  
 E-mail: [fmarozzo@deis.unical.it](mailto:fmarozzo@deis.unical.it),  
[trunfio@deis.unical.it](mailto:trunfio@deis.unical.it)

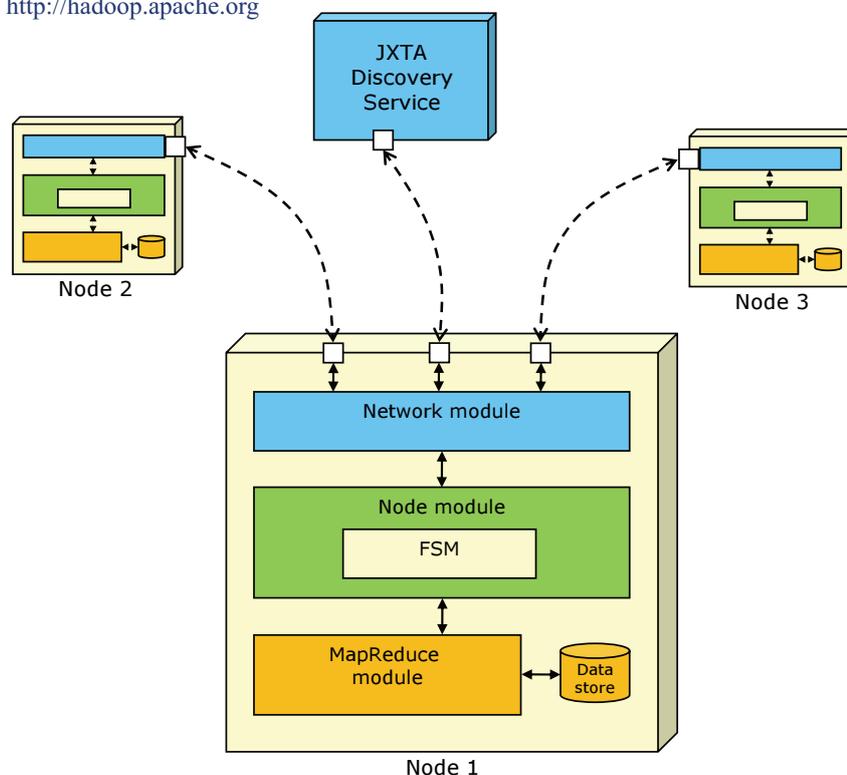


Figure 2: Software architecture of the P2P-MapReduce framework. Each node includes three software modules/layers: Network, Node and MapReduce. The Network module provides communication mechanisms with the other nodes and with the JXTA Discovery Service. The Node module implements the logic of the finite state machine shown in Figure 1. The MapReduce module manages the local execution of jobs and tasks.

# Considering Data Locality for Parallel Video Processing

by Rainer Schmidt and Matthias Rella

*Researchers at the Austrian Institute of Technology (AIT) are exploring ways to utilize cloud technology for the processing of large media archives. The work is motivated by a strong demand for scalable methods that support the processing of media content such as can be found in archives of broadcasting or memory institutions.*

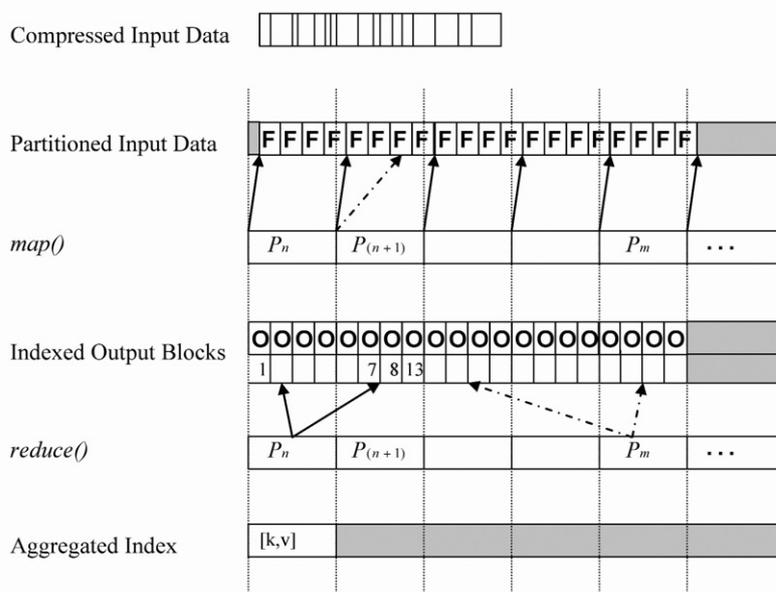
Infrastructure as a Service (IaaS) is a resource provisioning model that enables customers to access large-scale computer infrastructures via services over the Internet. It allows users to remotely host data and deploy individual applications using resources that are leased from infrastructure providers. A major strength of this approach is its broad applicability which is supported by the separation of application implementation and hosting. One of the most prominent frameworks that utilizes the IaaS paradigm for data-intensive computations has been introduced by Google. MapReduce (MR) implements a simple but powerful programming model for the processing of large data sets that can be executed on clusters of commodity computers. The framework targets applications that process large amounts of textual data (as required, for example, when generating a search index), which are parallelized on a master-worker principle. Scalability and robustness are supported through features like distributed and redundant

storage, automated load-balancing, and data locality awareness.

Here, we describe a method that exploits MapReduce as the underlying programming model for the processing of large video files. An application has been implemented based on Apache Hadoop, which provides an open-source software framework for data-intensive computing that can be deployed on IaaS resources. Typical use cases for the processing of archived video materials are for example file format migration, error detection, or pattern recognition. Employing a data-intensive execution platform for video content is desirable in order to cope with the large data volumes and the complexity introduced by diverse file and encoding formats. In previous work, we have developed a service that provides access to clusters of virtualized nodes for processing a large number of relatively small files like documents and images. In this application, parallelization takes place on a

per-file basis and the workload is decomposed into a list of file references. During execution, a worker node processes one file for each task, which is retrieved from a shared storage resource. Due to the nature of this problem, the application achieved reasonable speedup when executed within a cluster. Significant IO overhead, however, is introduced by the required file transfer between the compute nodes and the storage service. This necessitates the employment of a strategy that exploits data locality for large data volumes.

File systems like GFS or HDFS are designed to store large amounts of data across a number of physical machines. Files are split into chunks and distributed over the local storage devices of the cluster nodes. This allows the cluster management software to exploit data locality by scheduling tasks closely to the stored content (for example on the same machine or rack). Hence, worker nodes are preferentially assigned to process data that resides within a local



*Figure 1: A Distributed MR Video Processing Algorithm.*

partition. The approach has been proven to scale well for the processing of large text files. An interesting question is to explore how it can be applied to binary content as well.

One may consider pattern recognition in video files as a data-intensive task. As an example, we have implemented a MapReduce application, which (1) takes a video file as input, (2) executes a face recognition algorithm against the content, and (3) produces an output video that highlights the detected areas.

Video file formats typically provide a container that wraps compressed media tracks like video and audio, as well as metadata. For the pattern matching application, we solely consider the video track, which must be placed on the distributed file system for further processing. This has been done using a custom file format, which creates a raw

data stream providing a constant bitrate per frame, as shown in Figure 1. The raw data is automatically partitioned in blocks (64MB) by the file system and dispersed over the physical nodes P. During the first execution phase (map), the cluster nodes are assigned with data portions that correspond to single video frames. In this step, the input frames F are read and pattern recognition takes place. For each frame F, a corresponding output frame O is written and registered within a local index. This index provides key-value pairs that translate between a frame identifier and a pointer to the corresponding data. In an ideal case, each cluster node will only read and write from/to the local data partition. Indexing, however, is required as the data source and order of the frames processed by a cluster node are subject to load balancing and cannot be determined in advance. Also, a minimal data transfer across partitions is required for frames that are split

between data partitions. During the second processing phase (reduce), the locally generated index maps are reduced into a single result map. Support is provided by a corresponding driver implementation that directly retrieves the output video stream from the distributed file system based on its index map.

The application has proved to scale well in different test settings. Even small clusters of 3-5 computer nodes have shown performance improvements of up to 50% compared to the execution time of a sequential application.

**Link:**

<http://dme.ait.ac.at/dp>

**Please contact:**

Rainer Schmidt

Austrian Institute of Technology /  
AARIT

E-mail: [rainer.schmidt@ait.ac.at](mailto:rainer.schmidt@ait.ac.at)

## Online Gaming in the Cloud

by Radu Prodan, Vlad Nae and Thomas Fahringer

**Computational Grids and Clouds remain highly specialized technologies that are only used by scientists and large commercial organizations. To overcome this gap, University of Innsbruck is conducting basic research that is unusual compared with previous academic research projects in that it addresses a new class of application that appeals to the public for leisure reasons: Massively Multiplayer Online Games. Online games have the potential to raise strong interest, providing societal benefits through increased technological awareness and engagement. By standardizing on a Cloud-based platform and removing the need for large investments in hosting facilities, this research may remove the technical barrier and the costs of hosting MMOGs, and thus significantly increase the number of players while keeping the high quality of responsiveness of action games.**

Online entertainment, including gaming, is a strongly growing sector worldwide. Massively Multiplayer Online Games (MMOG) grew from 10,000 subscribers in 1997 to 6.7 million in 2003 and the rate is accelerating, with the number of subscribers estimated to be 60 million by 2011.

Today, MMOGs operate as client-server applications, in which the game servers simulate a persistent world within a game session, receive and process commands from the players distributed in the Internet (shootings, collection of items, chat), and interoperate with a billing and accounting system. Game servers are typically hosted by specialized companies called Hosters that rent to game operators computational and network capabilities for running game

servers with guaranteed Quality of Service (QoS).

To support millions of active concurrent players and many other entities simultaneously, Hosters install and operate a large static infrastructure, with hundreds to thousands of computers onto which the load of each game session is distributed. However, the demand of a MMOG is highly dynamic and depends on various factors such as game popularity, content updates, or weekend and public holiday effects. To sustain such highly variable loads, game operators over-provision a large static infrastructure capable of sustaining the game peak load, even though a large portion of the resources is unused most of the time. This inefficient resource utilization has negative economic impacts by

preventing any but the largest hosting centres from joining the market, and dramatically increases prices.

Today, acquiring own parallel computers is becoming less and less attractive to application developers or operators, since this is usually constrained by budget limitations, requires high-operational costs, and is ultimately affected by hardware deprecation following the Moore's law. To address this problem, a new research direction, known as Cloud computing, proposes an alternative by which resources are no longer hosted by the researchers' computational facilities, but leased from large specialized data centres only when and for how long they are needed. This frees institutions from permanent maintenance costs and eliminates the burden of hard-

ware deprecation. Through a new concept of “scaling-by-credit-card”, Clouds promise to immediately scale up/down an infrastructure according to the temporal needs in a cost-effective fashion. Moreover, the concept of hardware virtualization can represent a significant breakthrough for automating the deployment process of complex software that today remains a tedious and manual process that requires manual intervention of skilful computer scientists. Finally, the provisioning of resources through business relationships constrains specialized data centre companies in offering a certain degree of QoS encapsulated in Service Level Agreements (SLA) that significantly increases the reliability and the fulfilment of user expectations.

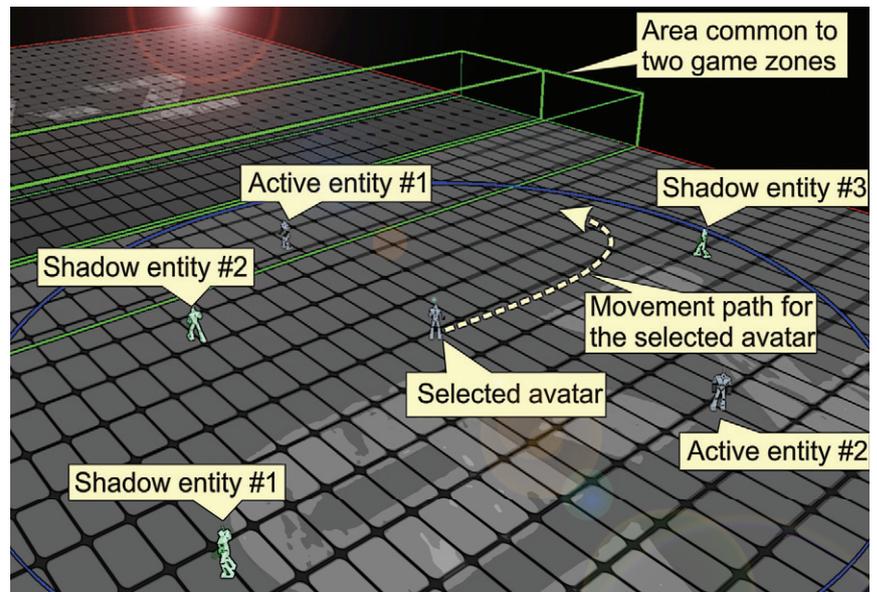
Despite the existence of many vendors that, similar to Grid computing, aggregate a potentially unbounded number of resources, Cloud computing remains a domain dominated by Web hosting or data-intensive applications, and whose suitability for computationally-intensive applications remains largely unexplored.

At the University of Innsbruck we are conducting basic research that aims to augment existing Cloud technologies with generic methods for QoS provisioning for real-time computationally-intensive applications. The ultimate goal is to apply and validate the generic research methods to MMOGs, as a novel class of socially important applications with severe real-time computational requirements to achieve three innovative technical objectives:

- Improved scalability of a game session hosted on a distributed Cloud infrastructure to a larger number of online users than the current state-of-the-art (ie 64-256 for First Person Shooter action games);
- Cheaper on-demand provisioning of Cloud resources to game sessions based on exhibited load;
- QoS enforcement with seamless load balancing and transparent migration of players from overloaded servers to underutilized ones within and across different Cloud provider resources.

These important goals are being technically achieved by investigating:

- Performance models for virtualized Cloud resources, including character-



*Figure 1: Snapshot from a First Person Shooter action game demonstrator developed at the University of Münster, Germany. The selected avatar represents the player that took the snapshot. The session is parallelized through game world, zoning on a large amount of servers, from which two servers with a common area for hiding avatar migration latencies is displayed for readability reasons. The two zones are further parallelized using entity replication: some of the entities close to the player are managed by the same server (the white active entities), while others are managed by another server (the light-green shadow entities).*

ization of the Cloud virtualisation and software deployment overheads;

- Proactive dynamic scheduling strategies based on QoS negotiation, monitoring, and enforcement techniques;
- SLA provisioning models based on an optimized balance between risks, rewards, and penalties;
- Resource provisioning methods based on time/space/cost renting policies, including comparative analysis between Cloud resource renting and conventional parallel/Grid resource operation.

This research originally started as part of the IST-034601 edutain@grid project which successfully completed in September 2009, and is currently being continued as part of a national project funded by the Austrian Science Fund (TRP-72-N23).

As part of these research activities, the University of Innsbruck is strongly cooperating with the Delft University of Technology in the areas of Cloud benchmarking (within the CoreGRID Virtual environments topic) and resource management (within the Scheduling topic), and with the University of Münster in the areas of MMOG load balancing and scalability (within the Service Level Agreements topic).

#### Links:

<http://www.edutaingrid.eu/>  
<http://www.dps.uibk.ac.at/mmog/>

#### Please contact:

Radu Prodan  
 University of Innsbruck, Austria  
 Tel: +43 512 507 6445  
 E-mail: radu@dps.uibk.ac.at

# Mastering Data-Intensive Collaboration and Decision Making through a Cloud Infrastructure

by Nikos Karacapilidis, Stefan Rüping and Isabel Drost

*Collaboration and decision making settings are often associated with huge, ever-increasing amounts of multiple types of data, obtained from diverse sources, which often have a low signal-to-noise ratio for addressing the problem at hand. In many cases, the raw information is so overwhelming that stakeholders are often at a loss to know even where to begin to make sense of it. In addition, these data may vary in terms of subjectivity and importance, ranging from individual opinions and estimations to broadly accepted practices and indisputable measurements and scientific results. Their types can be of diverse level as far as human understanding and machine interpretation are concerned.*

Nowadays, big volumes of data can be effortlessly added to a database; the problems start when we want to consider and exploit the accumulated data, which may have been collected over a few weeks or months, and meaningfully analyse them with the goal of making a decision. There is no doubt that when things get complex, we need to identify, understand and exploit data patterns; we need to aggregate big volumes of data from multiple sources, and then mine it for insights that would never emerge from manual inspection or analysis of any single data source.

Taking the above issues into account, the recently funded Dicode project (FP7-ICT-2009-5) aims to facilitate and augment collaboration and decision making

in data-intensive and cognitively-complex settings. To do so, it will exploit and build on the most prominent high-performance computing paradigms and large data processing technologies - such as cloud computing, MapReduce, Hadoop, Mahout, and column databases - to meaningfully search, analyse and aggregate data existing in diverse, extremely large, and rapidly evolving sources. Services to be developed and integrated in the context of the Dicode project will be released under an open source license.

The Dicode project is timely for the following reasons:

- Cloud computing is making a growing presence in both industry and academia. It is becoming a scalable serv-

ices delivery and consumption platform for Services Computing (at the same time, services are becoming more and more data intensive). Compared to its predecessors (ie grid computing, utility computing), cloud computing is better positioned in terms of economic viability, cost-effectiveness, scalability, reliability, interoperability, and open source implementations.

- There is much advancement in the development of scalable data mining frameworks and technologies (most of them exploiting the cloud computing paradigm), such as MapReduce, Hadoop, and Mahout. Likewise, text mining technologies (such as named entity recognition, named entity disambiguation, relation extraction, and

opinion mining) have reached a level in which it is - for the first time - practically feasible to apply semantic technologies to very large data collections, thus allowing capture of an unprecedented amount of information from unstructured texts.

- In parallel, there is much advancement in the development and decision making support applications, mainly by exploiting the Web 2.0 features and technologies.

• While helpful in particular problem instances and settings, the above categories of advancements demonstrate a series of limitations and inefficiencies

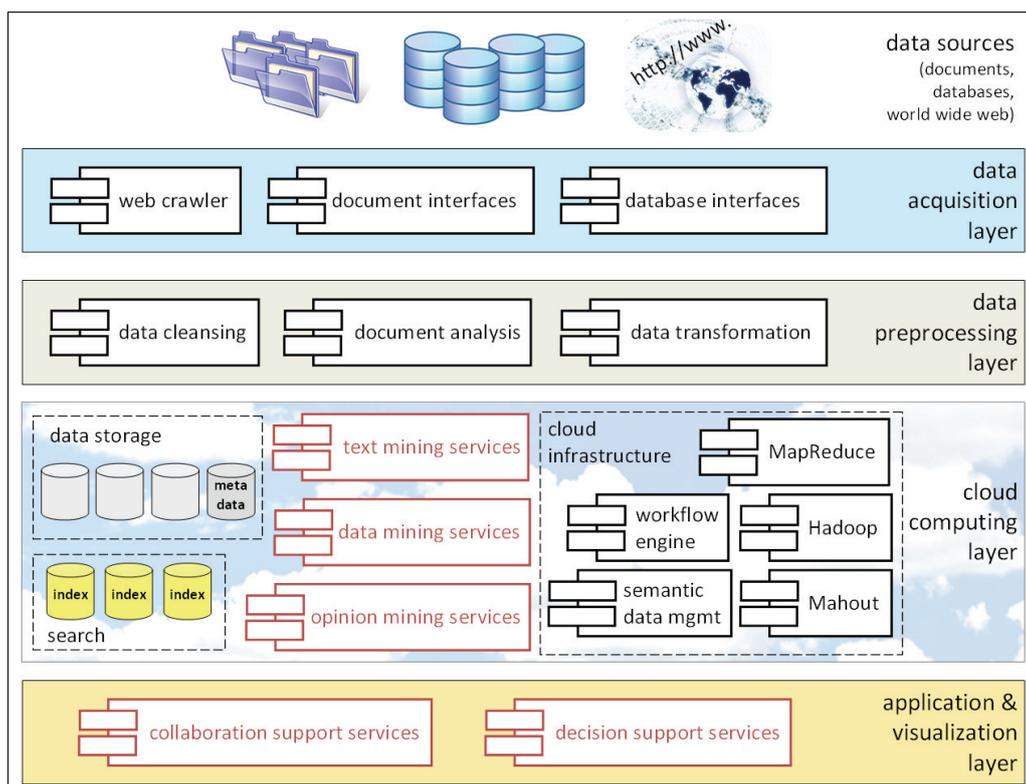


Figure 1: The Dicode Architecture and Suite of Services.

when applied to data-intensive and cognitively-complex collaboration and decision making support settings.

Building on current advancements, the solution foreseen in the Dicode project will bring together the reasoning capabilities of both the machine and the humans. It can be viewed as an innovative “workbench” incorporating and orchestrating a set of interoperable services (see Figure 1) that reduce the data-intensiveness and complexity overload at critical decision points to a manageable level, thus permitting stakeholders to be more productive and concentrate on creative and innovative activities. The achievement of the Dicode project’s goal will be validated through three use cases. These were chosen to test the transferability of Dicode solutions in different collaboration and decision making settings, associated with diverse types of data and data sources, thus covering the full range of the foreseen solution’s features and functionalities. In brief, these cases concern:

- Clinico-Genomic Research Assimilator. This case will demonstrate how Dicode can support clinico-genomic scientific research in the current post-genomic era. The need to collaboratively explore, evaluate, disseminate and diffuse relative scientific findings and results is more than profound today. To this end, Dicode envisages

the development of an integrated clinico-genomic knowledge discovery and decision making use case that targets the identification and validation of predictive clinico-genomic models and biomarkers. The use case is founded on the seamless integration of both heterogeneous clinico-genomic data sources and advanced analytical techniques provided by Dicode.

- Trial of Rheumatoid Arthritis Treatment. This case will benefit from Dicode’s services to deliver pertinent information to communities of doctors and patients in the domain of Rheumatoid Arthritis (RA). RA treatment trials will be carried out by an academic research establishment on behalf of pharmaceutical company. Each trial will evaluate the effectiveness of treatment for RA by analysing the condition in wrists (and possibly other joints). Dicode services will be used to enable an affective and collaborative way of working towards decision making by various individuals involved (Radiographers, Radiologists, Clinicians, etc.).
- Opinion Mining from unstructured Web 2.0 data. It is paramount today that companies know what is being said about their services or products. With the current tools, finding who and what is being said is literally searching for a needle in the haystack of unstructured information. Through

this case, we aim to validate the Dicode services for the automatic analyses of this voluminous amount of unstructured information. Data for this case will be primarily obtained from spidering the Web (blogs, forums, and news). We will also make use of different APIs from various Web 2.0 platforms, such as microblogging platforms (Twitter), and social network platforms (Facebook).

**Link:**

<http://dicode-project.eu/>

**Please contact:**

Nikos Karacapilidis  
Research Academic Computer  
Technology Institute, Greece  
Tel: +30 2610 960305  
E-mail: karacap@cti.gr

Stefan Rüping  
Fraunhofer Institute for Intelligent  
Analysis and Information Systems,  
Germany  
Tel: +49 2241 143512  
E-mail:  
stefan.rueping@iais.fraunhofer.de

Isabel Drost  
neofonie GmbH, Germany  
E-mail: isabel@neofonie.de

## ComCert: Automated Certification of Cloud-based Business Processes

by Rafael Accorsi and Lutz Lowis

***A key obstacle to the development of large-scale, reliable Cloud Computing is the difficulty of timely compliance certification of business processes operating in rapidly changing Clouds. Standard audit procedures are hard to conduct for Cloud-based processes. ComCert is a novel, well-founded approach to enable automatic compliance certification of business process with regulatory requirements.***

Reliable Cloud Computing must provide control over business process compliance. The central task to this end is certifying business processes for their adherence to regulations. However, due to the dynamics of Clouds, current manual audits for compliance certification, such as SAS-70 or SAS-117, are hard to apply in this setting. This is at odds with the increased flexibility that Clouds offer and with which companies can adapt

their business processes on demand. Consequently the lack of automated audit methods and the resulting risks of noncompliance currently inhibit enterprises from outsourcing their tasks onto the Cloud and prevent the full realization of the economic potential of Cloud Computing.

ComCert is a method of automated compliance certification of business

processes. Intuitively, auditors use ComCert tool support to check Cloud-based processes for adherence to a large set of different compliance requirements. The analysis carried out by ComCert is able to detect vulnerabilities arising from the data flow and control flow perspectives, eg, whether all required activities are included and whether activities happen in the prescribed order. If the process is com-

pliant, the approach generates evidence of correctness. If not, it provides counterexamples that identify violations of the compliance policies and indicate the vulnerable spots in the process. In doing so, ComCert complements and extends the research on business process verification, which has traditionally focused on checking the compatibility of communicating processes, ie absence of deadlocks and guarantee of service.

ComCert employs Petri nets as a formal basis to decide on the policy adherence of business processes. Petri nets provide an expressive, notation-independent formalism to capture the semantics of business processes. Also, compliance requirements can be expressed as usage

Process Modelling Notation and Business Process Execution Language into Petri nets exist, ComCert does not rely on a particular process notation for analysis, thereby being equally suitable for different Cloud-providers in a non-invasive manner. The use of Petri nets for the formalization of compliance requirements allows the circumscription of policy patterns. The refinement from regulations to policies can thus be captured in a formal and unambiguous, yet easily accessible way.

Specifically, ComCert builds upon an extensive classification of compliance requirements drawn from major regulations such as the Sarbanes-Oxley Act, the Health Insurance Portability and

each other, the conflicts are flagged, eg, if a regulation requires the retention of data while a competing regulation demands its deletion.

It is important to note that the compliance certificates regard only the process model and not the organizational and technical layers. General policies such as creating backups and logs, securing the network, raising the workforce's security awareness or having a security management process with risk assessments are an integral part of security efforts, but they do not pertain to a specific process model. Regarding the certification of the system components and procedures, for example, the execution engine and virtualization issues, additional vulnerability analysis and standard procedures as from COBIT and ISO 17799 would need to be considered.

Case studies using ComCert show that the "push-button" compliance check is feasible for industrial business processes. Ongoing investigation focuses on extending the kinds of analysis that can be carried out with ComCert. In particular, we are experimenting with information flow analysis and control techniques for Petri nets to detect interferences between communicating processes. These interferences denote further, more subtle but equally threatening design vulnerabilities that have not yet been considered in the certification of Cloud-based business processes.

Taking stock, ComCert contributes to automating the compliance certification of processes in the Cloud and consequently, fosters wider Cloud deployment and the compliant implementation of Cloud-based business models.

**Link:**  
<http://www.telematik.uni-freiburg.de/comcert>

**Please contact:**  
 Rafael Accorsi and Lutz Lowis  
 Dept. of Telematics, University of Freiburg, Germany  
 Tel: +49 761 203 4926  
 E-mail: [accorsi@iig.uni-freiburg.de](mailto:accorsi@iig.uni-freiburg.de),  
[lowis@iig.uni-freiburg.de](mailto:lowis@iig.uni-freiburg.de)

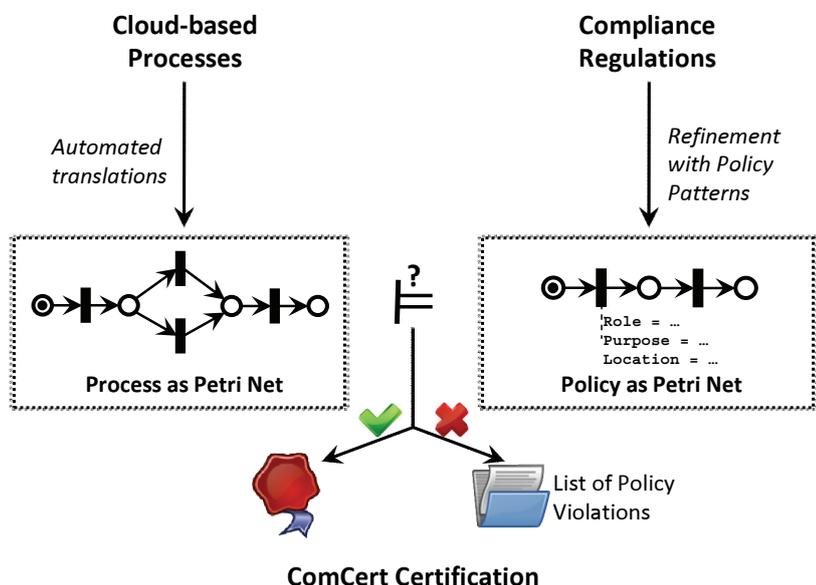


Figure 1: The ComCert approach of checking processes for compliance.

control policies, and Petri nets serve as formal representation of those policies. Given the Petri net representation of both a process and the applicable policies, the compliance check through ComCert certification is reduced to a type of reachability problem in Petri nets. Put simply, the goal is to demonstrate that the process modelled in one Petri net (the "process net") satisfies the compliance requirements modelled in another Petri net (the "policy net"). Petri net reachability is a well-investigated problem for which efficient algorithms and, hence, tool support for automation exists.

Since translations from standard process notations such as Business

Accountability Act and the PATRIOT Act. The resultant classification consists of three classes, where each requirement either (1) requires certain activities to (not) be performed before or after other activities, (2) describes the mandatory flow of data between activities and (3) prescribes additional conditions on data, eg pseudonymization and retention. Petri nets formalize these classes as high-level patterns, thereby reducing the effort for the specification of requirements. For a particular application domain and regulation, concrete requirements are stepwise refined into instances of high-level patterns. With instantiated patterns, it is also possible to detect policy inconsistencies in an automated manner. If the policy nets contradict

# Buiding Discrete Spacetimes by Simple Deterministic Computations

by Tommaso Bolognesi

*The Computational Universe conjecture relates complexity in physics with emergence in computation. Our current research efforts are meant to put the surprisingly powerful notion of (computational) emergence at the service of recent quantum gravity theories, with special attention to the Causal Set Programme, which assumes causality among events as the most fundamental structure of spacetime, and causal sets as the most appropriate way to describe it.*

Our physical universe is discrete, finite, unbounded, deterministic and computational. Of course most readers will disagree with most of these attributions, but a number of researchers in the last few decades have been willing to take at least some of them as stimulating ‘working hypotheses’ for exploring alternative physical theories whose basic ideas could hardly be beaten in terms of simplicity.

Discrete means that there exists a tiniest scale at which the fabric of space (and also of space-time) appears as a pomegranate, made of indivisible atoms, or seeds; this viewpoint is adopted in theories such Loop Quantum Gravity and in the so called Causal Sets Programme, and is reflected in models such as Penrose’s spin networks and spin foams.

*Finite* means that the number of seeds in the pomegranate is finite, say  $10^{234}$  (as of year 2010 a.c.)

*Unbounded* means that new seeds keep popping up as the universe evolves, but always in finite quantities, perhaps one at a time.

*Deterministic* means that at this ultimate level, reality obeys precise rules that do not involve any coin flipping; we are thus assuming that God indeed does not play dice, and we look with great hope at some recent efforts, eg by G. ‘tHooft, that try to unveil a deterministic layer under the apparent probabilistic nature of Quantum Mechanics.

Computational means that these rules can be implemented and executed step by step on a digital computer; this does not mean that we have to postulate the existence of a divine digital Computer that sits in some outer space and runs the program for our universe, for the same reason that under a continuous mathematics viewpoint we do not need to postulate the existence of a divine analog Computer that runs the Navier-Stokes differential equations of fluid dynamics.

Eminent physicist Richard Feynman is one of the scientists that have been intrigued by the idea of a discrete, computational universe. The following is a famous passage from his



1964 Cornell Lectures (‘The Character of Physical Law’): ‘It always bothers me that, according to the laws as we understand them today, it takes a computing machine an infinite number of logical operations to figure out what goes on in no matter how tiny a region of space, and no matter how tiny a region of time. How can all that be going on in that tiny space? Why should it take an infinite amount of logic to figure out what one tiny piece of spacetime is going to do? So I have often made the hypothesis that ultimately physics will not require a mathematical statement, that in the end the machinery will be revealed, and the laws will turn out to be simple, like the chequer board with all its apparent complexities.’

What type of complexity can appear, or emerge, by playing a simple game on a checkerboard? An answer is provided by Conway’s Game of Life, a well known example of two-dimensional cellular automaton that became popular in the 1970’s. By letting all elements of a square array of cells obey, synchronously, the same simple rule, which only refers to the color (black or white) of the cell and of its eight neighbors, one obtains surprisingly complex populations of moving patterns (‘gliders’, kites’, ‘darts’, ...) that suggest a lively aerial scenario.

The fact that simple deterministic computational rules can originate highly complex patterns and dynamics has been further explored and popularized by Stephen Wolfram, who showed that even simpler models of computation, notably one-dimensional, two-color cellular automata (ECA), can exhibit very complex behaviours. Two surprising examples are provided by ECA n. 30, with its pseudo-random computations, and ECA n. 110, with its emergent particles. The latter is illustrated in Figure 1, where the configurations of the 1-D array of cells are stacked, so that the horizontal and vertical dimensions correspond, respectively, to space and time. While these particles and their behaviors are clearly

reminiscent of scattering in 'real' physical phenomena, Cook and Wolfram were able to formally prove that they can also simulate any Turing machine (like Conway's Game of Life), thus turning the device into a universal computer.

There are several ways in which one can represent the computations of a given model, devise complexity indicators characterizing their behaviours, and detect the emergence of interesting features such as pseudorandomness or interacting particles. A particularly attractive way to do this is to consider the 'causal set' associated with the computation, which is formed by a set of events and a set of causal relations among them. This approach appears as particularly appropriate for applications in fundamental physics, since causal sets are regarded as one of the most appropriate models of discrete, physical space-time.

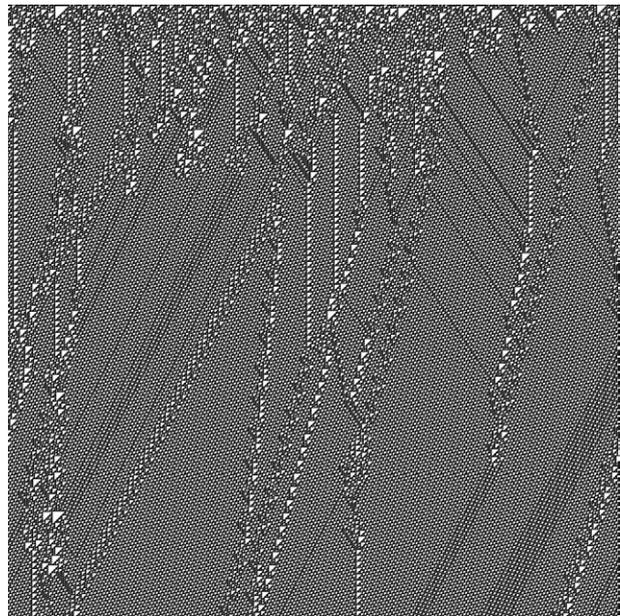
Obtaining causal sets from the computations of simple models such as Turing machines, network mobile automata, or graph rewrite systems, is fairly easy.

In doing so we may obtain two attractive results. On one hand, these 'algorithmic' causal sets might represent, under the 'Computational Universe perspective, the only information of physical relevance that we can extract from the considered models of computation, and a way to abstract from the internal machinery of the latter. On the other hand, this approach might represent a fully deterministic, radical alternative to the probabilistic techniques currently adopted in the Causal Set Programme for growing discrete spacetimes.

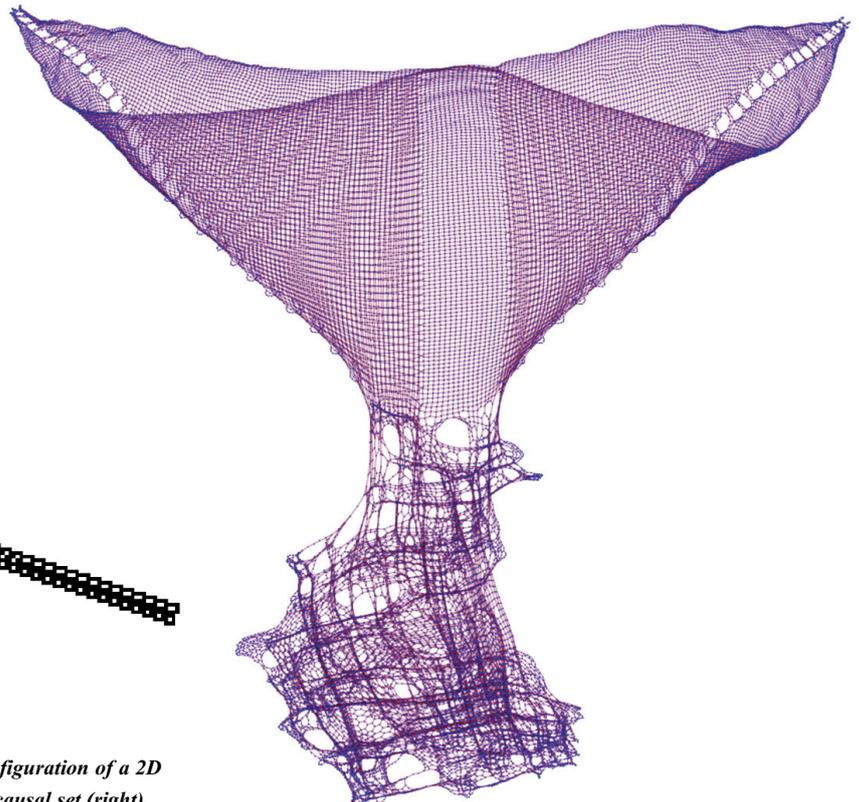
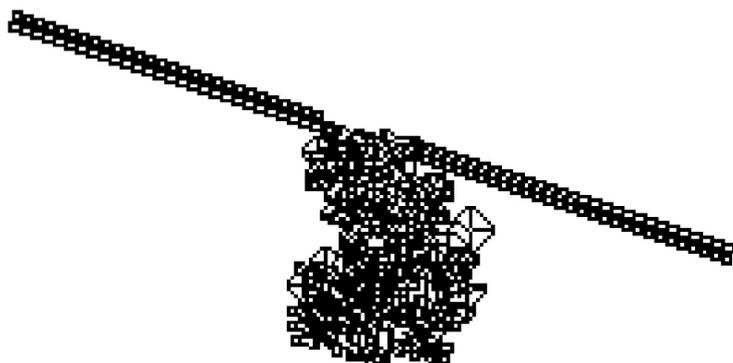
If pseudo-particles and pseudo-randomness emerge from simple deterministic computations, we may expect to spot similar phenomena also in the causal sets derived from them. An example of two coupled particles that we have discovered, emerging from the computation of a 2D Turing machine (one moving on a checkerboard), is shown in Figure 2.

**Link:**  
<http://arxiv.org/abs/1004.3128>

**Please contact:**  
Tommaso Bolognesi  
ISTI-CNR, Italy  
E-mail: [tommaso.bolognesi@isti.cnr.it](mailto:tommaso.bolognesi@isti.cnr.it)



*Figure 1: The artificial 'particles' emergent from the computations of Elementary Cellular Automaton n. 110 are reminiscent of 'real' scattering phenomena in physics, and at the same time achieve computational universality (Turing-completeness).*



*Figure 2: Emergence of two coupled particles in the final configuration of a 2D Turing machine computation (left) and in the corresponding causal set (right).*

# Improving the Security of Infrastructure Software using Coccinelle

by Julia Lawall, René Rydhof Hansen, Nicolas Palix and Gilles Muller

*Finding and fixing programming errors in deployed software is often a slow, painstaking, and expensive process. In order to minimise this problem, static analysis is increasingly being adopted as a way to find programming errors before the software application is released. Coccinelle is a program matching and transformation tool that makes it easy for developers to express static analysis-based software defect-finding rules and scan software source code for potential defects.*

Defects are continually found in end-user software, even in seemingly reliable software on which millions of people depend. Finding and fixing defects is slow and painstaking: a tester or user finds that the software crashes, that person submits a bug report, a maintainer studies the code to find the root cause of the reported problem, and finally some change is made to fix it. The process then repeats when the next crash occurs. An alternative is to use static analysis, which is being increasingly adopted in both commercial and research tools. In this approach, a tool scans the software source code according to a collection of rules, and signals apparent defects. Static analysis is pervasive: all of the source code is checked, even source code that is rarely executed. Furthermore, defects are often reported at or near the place where the code needs to be changed; it is not necessary to connect an external run-time behavior to a particular source code element.

A static analysis approach, however, is only as good as the set of rules that it checks. Indeed, while the standard defect detection method relies on observed run-time problems to

initiate the defect-finding process, static analysis requires that a rule developer anticipate causes, and ideally solutions, of potential problems. What is needed is an approach to easily turn a maintainer's intuition about a potential problem, obtained from either debugging a runtime error or from studying the reports from a static analysis tool, into a rule that can be used to scan the software for similar problems.

During the past four years, supported in part by the Danish Research Council for Technology and Production at the Universities of Copenhagen and Aalborg in Denmark and the French National Research Agency at INRIA in France, we have been developing the Coccinelle program matching and transformation system. Coccinelle allows software developers to express rules as patterns described in terms of source code elements that can be abstracted such that they match not just one specific piece of code, but also similar code structures throughout a software project. Coccinelle patterns can furthermore be annotated using '-' and '+', following the commonly used patch syntax, to indicate code to remove and add, respectively. Thus, Coccinelle rules may not only find defects, but can also fix them. We have used Coccinelle to help in finding and fixing hundreds of defects in the Linux operating system. A typical example of a rule specification used in the context of Linux is shown in the figure. Coccinelle has furthermore been adopted by a number of developers outside of our research group, for use on both open and closed source software projects.

Our most recent work has lead in two new directions to better support the defect-finding process. First, to improve the robustness of a software system, it is necessary to understand how defects are introduced. While Coccinelle enables finding defects in multiple versions of a software project, it does not distinguish between defects that are long lasting and those that are continually fixed and reappear. To obtain a more complete picture, we have developed Herodotos, which correlates defect reports across multiple versions, even in the presence of other changes in the source code. A second direction is the integration of Coccinelle and Clang, the C language frontend developed as part of the LLVM (Low Level

Virtual Machine) compiler framework. By integrating Clang, Coccinelle is able to leverage the comprehensive program analysis capabilities found in both Clang and the underlying LLVM framework. This results in better precision for Coccinelle searches, and enables supporting a wider range of searches.

Coccinelle is open source and can be downloaded freely from the Coccinelle web page.

## Links:

Coccinelle: <http://coccinelle.lip6.fr/>  
 Herodotos: <http://coccinelle.lip6.fr/herodotos.html>  
 Low Level Virtual Machine (LLVM):  
<http://www.llvm.org/>  
 Clang: <http://clang.llvm.org/>

## Please contact:

Julia L. Lawall  
 University of Copenhagen / DANAIM, Denmark  
 Tel: +45 35321405  
 E-mail: [julia@diku.dk](mailto:julia@diku.dk)

```

diff -u -p a/drivers/dma/amba-p108x.c b/drivers/dma/amba-p108x.c
--- a/drivers/dma/amba-p108x.c 2010-08-20 01:26:52.000000000 +0200
+++ b/drivers/dma/amba-p108x.c 2010-08-24 18:19:25.000000000 +0200
@@ -616,7 +616,7 @@ static inline u32 p108x_pre_boundary(u32
     static int p108x_fill_llis_for_desc(struct p108x_driver_data *p108x,
                                       struct p108x_txd *txd)
     {
-         struct p108x_channel_data *cd = txd->cd;
+         struct p108x_channel_data *cd;
+         struct p108x_bus_data *mbus, *sbus;
         u32 remainder;
         int num_llis = 0;
@@ -630,6 +630,7 @@ static inline u32 p108x_pre_boundary(u32
     dev_err(&p108x->adev->dev, "%s no descriptor\n", __func__);
     return 0;
 }
--- mini_null_ref.out Top L1 (cocci)----6:34-----
@@
type T;
expression E;
identifier i,fld;
statement S;
@@
- T i = E->fld;
+ T i;
... when != E
   when != i
if (E == NULL) S
+ i = E->fld;
--- mini_null_ref.cocci All L1 (cocci)----6:34-----

```

Figure 1: Rule specification and matches in Coccinelle.

# Teaching Traffic Lights to Manage Themselves ... and Protect the Environment

by Dirk Helbing and Stefan Lämmer

*Adaptive techniques can modernize traffic control, saving fuel, reducing travel times and emissions, and doing it all without limiting drivers' mobility. This approach promises to save driving time and benefit the environment at the same time.*

Currently, traffic jams and road congestion do a lot more than annoy millions of people every day. In the United States alone, delays linked to backed-up traffic cost nearly \$100 billion each year, and waste more than 10 billion litres of fuel, not to mention countless human hours. Additionally, CO<sub>2</sub> and other pollutants are spewed into the atmosphere from traffic stalled in jams.

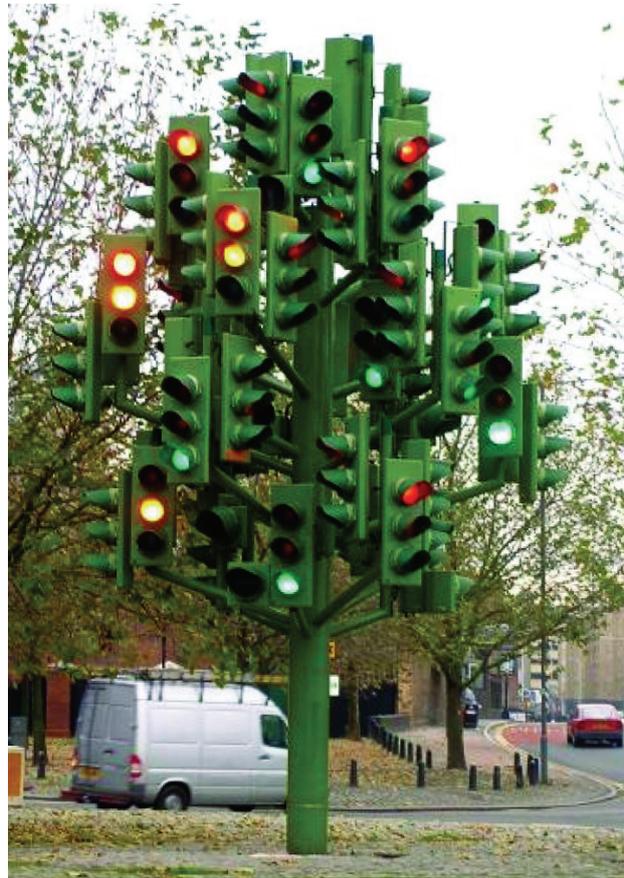
The new approach is based on giving traffic lights some simple traffic-responsive operating rules and letting the lights organise their own on-off schedules. Traffic is modelled as if it were a fluid, where traffic leaving one road has to enter another, like fluid moving through a network of pipes.

Jams can arise, obviously, if traffic entering a road overloads its capacity. To avoid this, each set of lights is given sensors that feed information about the traffic conditions at a given moment into a computer chip, which then calculates the flow of vehicles expected in the near future. The chip also calculates how long the lights should stay green in order to clear the road and thereby relieve the pressure. In this way, each set of lights can estimate for itself how best to adapt to the conditions expected at the next moment.

Simulations showed, however, that this simple rule isn't enough: the lights sometimes adapt too much. If they are only adapting to conditions locally, they might stay green for too long and cause trouble further away. The algorithm has been modified so that what happens at one set of traffic lights affects how the others respond. By working together and monitoring the lengths of queues along a stretch of road, the self-organised lights prevent long jams from forming.

Despite the simplicity of the rules, they seem to work remarkably well. Computer simulations demonstrate that lights operating this way would achieve a significant reduction in overall travel times and keep no one waiting at a light too long. One of the biggest surprises, however, is that all this improvement comes with the lights going on and off in an unpredictable way, not following a regular pattern as one might expect.

The adaptive control does not fight the natural fluctuations in the traffic flow by trying to impose a specified flow rhythm. Rather, it uses randomly appearing gaps in the flow to serve



*A 'special' traffic light in London.*

other traffic streams. According to simulations, the algorithm can reduce average delay times by 10%–30%. The variation in travel times goes down as well. Being responsive to local demands, when the traffic flow is low, an approaching car can be sensed and the light changes to green to let it through.

We are working with a German traffic agency to implement the idea. In tests based on Dresden's road layout, simulation results for realistic, measured traffic flow conditions have been very encouraging. The simulations show significant reductions in waiting times and fuel consumption.

Similar problems arise in many other systems including complex manufacturing plants, the supply chains of large organizations, the electricity grid, or the administrations of organizations. These will be studied in a large-scale technosocio-economic research project called FuturICT.

## **Links:**

<http://www.santafe.edu/media/workingpapers/10-09-019.pdf>

<http://www.patent->

[de.com/20100805/DE102005023742B4.html](http://www.patent-de.com/20100805/DE102005023742B4.html)

<http://futurict.eu>

## **Please contact:**

Dirk Helbing

ETH Zurich, Switzerland

E-mail: [dirk.helbing@gess.ethz.ch](mailto:dirk.helbing@gess.ethz.ch)

Stefan Lämmer

TU Dresden, Germany

E-mail: [traffic@stefanlaemmer.de](mailto:traffic@stefanlaemmer.de)



## A New Approach to the Planning Process makes Huge Savings for the Railway Sector

by Malin Forsgren and Martin Aronsson

*The new planning process, developed by SICS, is called Successive Allocation, and the general principle governing it is the separation of the train plan into two parts: (i) The service that Trafikverket (the Swedish Transport Administration) commits to deliver to its customers (the operators), and (ii) Production plans containing the details of how Trafikverket will deliver what has been promised. The latter is significantly improved by applying elements of lean production and just-in-time.*

Since 2005 SICS has been conducting timetable related research which has resulted in a prototype tool for efficient timetable generation. Just as important, research has also shown that there is huge potential to revise the entire planning process at Trafikverket (the Swedish Transport Administration), to enable seamless and continuous plan updates and to offer flexibility and better service to their customers.

The current projects are called Train Plan 2015 and Maraca, both stemming from the project The Dynamic Train Plan (2005-2008). While Maraca focuses on developing algorithms for an optimizing tool, Train Plan 2015 has been looking at (among other things) how the process at large needs to evolve in order to get the most benefit out of technological improvements like optimizing tools.

As of this year, the railway market in Sweden is completely deregulated, meaning that many operators now compete for the same capacity. It's the task of Trafikverket to coordinate the applications for train paths from all operators and create one single train plan for the whole country, for one full year at a time. The 2011 train plan takes shape between April and September 2010, takes effect in mid-December 2010 and is valid until mid-December 2011.

Establishing a very detailed plan as long as 15 months in advance (which is the case for the last trains in the plan) does not seem very efficient, particularly given the complexities involved with making changes to the plan. After examining the process, the researchers at SICS were able to show how applying elements of lean production and just-in-time would improve the planning process significantly.

The key is to minimize waste, and to "pull" rather than "push", throughout the process. Translated to our setting, lean production means that Trafikverket should avoid putting effort into producing anything (eg details in a plan) that is likely to become outdated before it will be used. The pull strategy involves producing only what the next step in the value chain really needs, when it is needed, and to perform only value-adding activities.

The new process is called Successive Allocation, and the general principle governing it is the separation of the train plan into two parts: (i) the service that Trafikverket commits to deliver to its customers (the operators), and (ii) production plans containing the details outlining how Trafikverket will deliver what has been promised. The first part is negotiated and clearly stated in contracts, and should typically include the arrivals and departures that are important to the operators. All other details belong in the production plans, and can change as often as necessary as long as the delivery commitments stay the same.

With the high capacity utilization of the railway network that we see today in Sweden, the change suggested by SICS could potentially result in huge savings for the railway sector, and for society as a whole, since it will make railway traffic more efficient and flexible. Trafikverket has set up a development program to have the Successive Allocation incorporated by the year 2015, and SICS will continue to play an important role while this change is being implemented.

**Link:**  
<http://www.sics.se/groups/railways>

**Please contact:**  
Martin Aronsson  
SICS; Sweden  
Tel: +46 8 633 15 87  
E-mail: [martin@sics.se](mailto:martin@sics.se)

# Fast Search in Distributed Video Archives

by Stephan Veigl, Hartwig Fronthaler and Bernhard Strobl

*Exciting perspectives are emerging in the field of visual surveillance. Due to the rapidly growing number of cameras and volume of video data, there is an increasing need for a method that enables quick pinpointing of significant data within the “sea” of irrelevance.*

Today’s visual surveillance systems reduce the input data to a great extent by simple motion detection; nevertheless, the resulting amount of data produced by such systems remains unmanageable. Therefore, further automated means for its analysis are required. For this purpose we resort to state-of-the-art visual surveillance algorithms for object detection, tracking and activity analysis. We developed a surveillance data analysis framework which is capable of efficiently assisting in data exploration from distributed video archives with different data formats and which allows rule and/or content-based object and event searches in large surveillance datasets.

With our system it is possible to seek a distinctive object in a huge archive of videos (multiple cameras, 24 hours). A typical use case is one or more surveillance cameras watching the entrance of a parking lot. We can for example, establish a rule to detect all cars entering or exiting the parking lot, but to ignore all vehicles just passing by. The system will detect all objects in the video and filter the results according to user defined rules. It will present a separate list for every rule with the matching objects. Additionally, we can search for a given sample image, resulting in a list, which is sorted by the according match score (see Figure 1).

## System Architecture

The design of our distributed high-performance archive search system defines the following components (see Figure 2):

- Multiple Video Archives / Camera Metadata Databases
- Analytics Core
- Configuration and Detection Database
- Several GUI Clients.

Each of the above services is supposed to run on a dedicated machine, optimized for the respective task. However, for demonstration purposes, it is also possible to run the whole system on a single computer.

The analytics core (see Figure 2) is a three stage system following a modular programming paradigm:

1. detection and tracking modules
2. filtering modules
3. matching modules.

At the moment we have implemented a blob tracking module (Moving Objects) and a person detection module in the first stage. As filtering module we use an Event Engine module as core of our rule based approach. A generic appearance-based matching module is used to sort the results in the matching stage. The above-mentioned modules are detailed in the following:

**Blob Tracker (Moving Objects):** This module does not specialize in any particular type of object, but rather detects every moving region (blob) in a scene. For this purpose, we employ a robust background model, which bases its decisions (foreground or background) on a compressed history of each pixel, referred to as codebook. For every new frame, each pixel is compared with its corresponding codebook and classified into either a foreground or a background pixel. This decision is based on the distribution of the previously observed pixel values.

**Person Detector:** Blob-based object detection suffers from sudden performance decay if the density of image objects becomes high and frequent dynamic occlusions are present. To overcome this problem, we have developed a human detection framework incorporating Bayesian statistics, which is capable of detecting and coarsely segmenting humans in moderately crowded scenes in real-time. Shape and motion cues are combined to obtain a maximum a posteriori solution for human configurations consisting

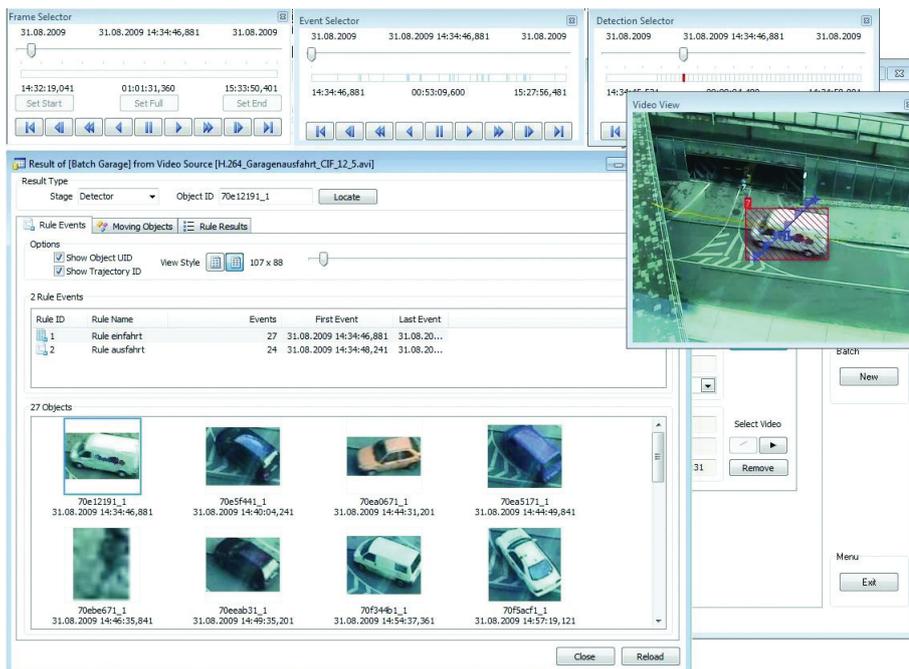


Figure 1: Detection results of user defined rule.

of many - possibly occluded - pedestrians viewed by a stationary camera.

High parallelization of the computations of both the blob tracker and the person detector enables an efficient implementation on graphics hardware. This yields real-time performance for the person detector and more than 100x real-time in case of blob tracking.

**Event Engine:** This is the central filtering module. The detected objects can be filtered by a highly flexible combination of user configurable events (eg crossing a tripwire or entering an area) and properties (eg object height or width). All rules on one video are processed in parallel. So, for instance, it is possible to statistically analyse the whole traffic-flow in a roundabout, or similar, using a single-pass video processing technique.

**Generic Appearance-Based Matcher:** An appearance-based matching functionality allows the user to search for an object

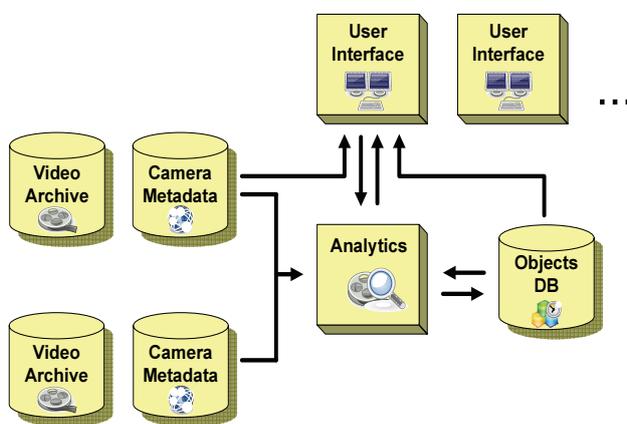


Figure 2: System overview.

with specific appearance attributes. The search is initiated by specifying a query image, which can be provided either by selecting an image region in a video or uploading a bitmap. An object descriptor is computed based on the covariance matrix of the image feature vector. This target descriptor is compared to the descriptors of found objects producing a ranking of the objects.

In the future, we plan to augment the analytics core (see Figure 2) with a face and a license plate detector together with the according matching modules.

**Please contact:**

Stephan Veigl  
AIT Austrian Institute of Technology GmbH  
Tel: +43 50550-4270  
E-mail: [stephan.veigl@ait.ac.at](mailto:stephan.veigl@ait.ac.at)

## Meeting Food Quality and Safety Requirements with Active and Intelligent Packaging Techniques

by Elisabeth Ilie-Zudor, Marcell Szathmári, Zsolt Kemény

*In numerous regions worldwide customers' expectations and legislations requirements regarding food quality and safety are changing. Customers now turn more and more towards fresh food with little intervention in its raw materials, while legislation regarding tracking of perishable products along the entire supply chain becomes more rigorous. These changes pose an increasing challenge both for the food and packaging industries.*

The usage of active and intelligent packaging (AIP) can help dealing with the new requirements for higher food quality and safety. Active packaging refers to solutions with a quality-preserving function for the contents without altering the composition of the product itself, while intelligent packaging covers the solutions with diagnostic and indicator functions for assessing an estimated shelf life or supply information on actual product freshness.

Nowadays, considerable development and application efforts are spent, especially in the USA, Japan and Australia, to integrate active and intelligent functions in high-end (consumer) packaging types for food and beverages. In Europe, the use of these innovative materials and technologies is limited to large and internationally operating food and packaging producers, and effective active and intelligent materials for food packaging applications are still rather rare on the market. Aside from less compatible legislative framework, the spreading of active and intelligent packaging solutions in Europe is hampered by a lack of knowledge, application results, or dedicated development of these new materials and technologies for food packaging applications. The research base is fragmented: information regarding effectiveness and the reliability of these types of active and intelligent systems to improve the food quality and safety in different food sectors is widespread, but not easy to detect and not always comparable. For a successful launch of active or intelligent packaging, the choice of food product in combination with the choice of the type of active and intelligent concept is important.

Information exchange regarding knowledge and promising technologies on active and intelligent packaging is especially needed for SMEs who are not in the position to do this analysis on their own.

Starting its development in September 2009, the "Development of tools to communicate advanced technologies on active and intelligent packaging to meet the needs and trends in food processing and retailing and to improve the knowledge transfer especially for SMEs" project has been initiated by 12 SME associations and academic insti-

tutes from seven European countries (Austria, Belgium, Germany, Hungary, the Netherlands, Slovenia and Spain). The project bears the acronym AIP-Competence Platform and will have a development period of two years. It is funded under the framework of CORNET 6th Joint Call for Transnational Collective Research (<http://www.cornet-era.net>).

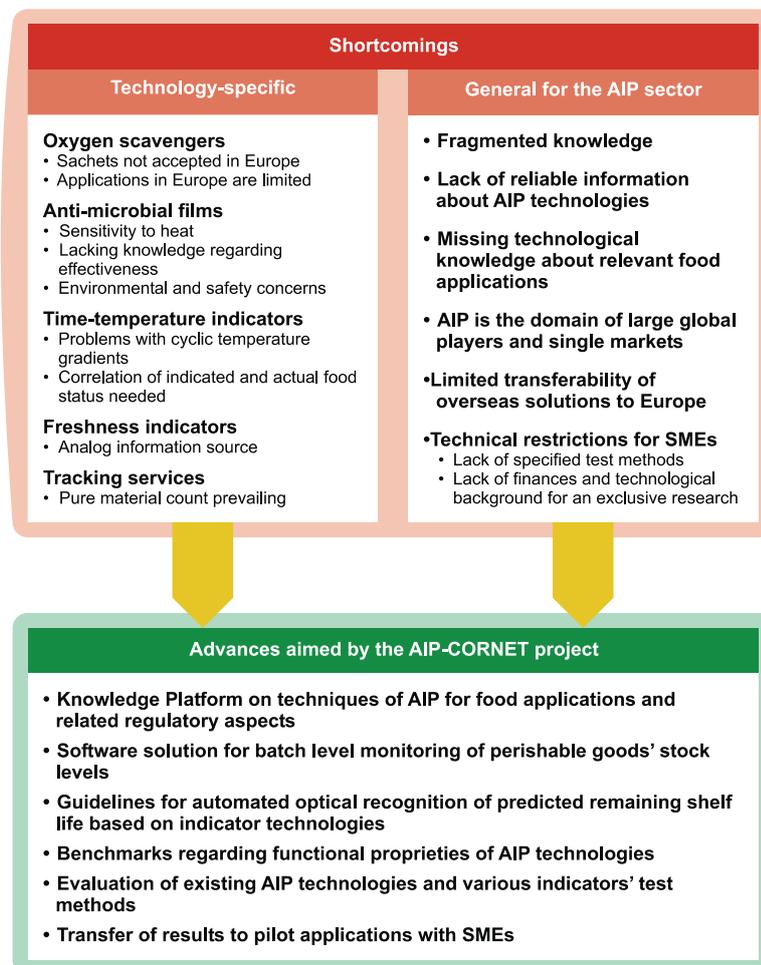
Due to the fact that the oxygen absorbers, the antimicrobial packaging materials, the time-temperature indicators, intelligent expiry date labels, supply chain monitoring via printable indicators and tracking information databases gathering all types of expiry and handling-related information along the supply chain are the most substantial market segments in active and intelligent packaging, the project will focus the work on these techniques.

The project will:

- collect technical knowledge and available data on commercialized techniques of active and intelligent packaging for food applications (technical data on working principles, active substances, functional properties, processing requirements, recommended application areas, suppliers of the active substances),
- review the available knowledge of active and intelligent packaging technologies and their application in the food supply-chain,
- report on the significance, accuracy of measurement and repeatability of existing test methods for characterizing oxygen scavenger systems, antimicrobial packaging materials, time-temperature-indicator systems, and freshness indicators,
- provide a software solution to monitor the stock levels of perishable goods on a batch level, providing also quality and stock-level alerts, in order to help companies prevent spoilage and better adjust their forecasts
- elaborate barcode symbology design guidelines for automated optical recognition of predicted remaining shelf life
- demonstrate and disseminate best practice reference samples on successful applications to extend shelf life and to monitor food quality and safety,
- deliver targeted information about impacts and minimum requirements of highly perishable foods to prolong the shelf life and to ensure the right quality and safety of the packed products, respectively up-to-date information on regulatory aspects,
- set up a group of key institutes and food and packaging specialists in Europe skilled in developing and testing of active and intelligent packaging as contact points for the food and packaging industry, especially SMEs, at national and EU level,
- provide a Knowledge and Communication Platform that can be leveraged to efficiently communicate the knowledge and data collected.

With all of the focal activities apparently centered around food chemistry and packaging technologies, the project still encompasses a number of IT-related developments as well, especially as far as the knowledge sharing platform and development of tailored trackable solutions are concerned.

As already addressed earlier, the technology transfer efforts of the project are to be supported by an online knowledge



*Problems addressed and goals pursued by the AIP-Competence Platform project.*

sharing platform, envisaged to serve both experts (large-scale producers, technology stakeholders, etc.) and newcomers (SMEs, small-scale producers, etc.) in the AIP knowledge domain. In order to serve the large diversity of needs represented by prospective visitors, the topic map paradigm is used as an organizing principle for a coherent corpus of resources, while different interfaces will support the wide range of search and browsing criteria of visitors.

Tracking and easy inclusion in tracking networks are another major challenge where information technology takes a notable share of the challenges. While the project consortium already has experience in SME-accessible tracking services at hand, specific solutions have to be elaborated for food supply chains. This also includes the development of uniquely identifiable intelligent packaging labels that present a low-cost alternative to, e.g., sensor-equipped RFID while remaining machine-readable. The elaboration of such lean solutions is expected to bridge gaps that, to date, deprive 'low-tech' participants of the benefits of transparent supply chains.

**Link:**

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

**Please contact:**

Elisabeth Ilie-Zudor

SZTAKI, Hungary

Tel: +36 1 279 6195

E-mail: [ilie@sztaki.hu](mailto:ilie@sztaki.hu)

## CLEF 2010: Innovation, Science, Experimentation

by Nicola Ferro

*Since 2000 the Cross-Language Evaluation Forum (CLEF) has played a successful role in stimulating research and promoting evaluation in a wide range of key areas in the information retrieval domain, becoming well-known in the international IR community. The results had traditionally been presented and discussed at annual workshops in conjunction with the European Conference for Digital Libraries.*



CLEF 2010 represented a radical innovation of this traditional format. The aim is to continue with the benchmarking evaluation experiments, but also to provide more space for analysis, reflection and discussion. It was thus organised as an four-day event: a peer-reviewed conference followed by a series of laboratories and workshops.

### CLEF 2010: The Conference

The conference aimed at advancing research into the evaluation of complex multimodal and multilingual information systems. Topics covered included experimental collections and datasets, evaluation methodologies, resources and tools. Two keynote talks by Norbert Fuhr, University of Duisburg-Essen and Ricardo Baeza-Yates, Yahoo! Research and Universitat Pompeu Fabra discussed future directions for experimental evaluation from academic and industrial perspectives. Reports from other major evaluation initiatives were presented: the Text REtrieval Conference, USA, the NII-NACSIS Test Collection for IR Systems, Japan, the INitiative for the Evaluation of XML Retrieval, Australia, the Information Retrieval Evaluation Seminar, Russia, and the Forum for Information Retrieval Evaluation in India. There were also two panels. The first presented the objectives of a new EU FP7 Network of Excellence, PROMISE, which aims at advancing the experimental evaluation of complex multimedia and multilingual information systems. In the second, researchers

responsible for creating and running some of the main evaluation initiatives of the last two decades, discussed what has been achieved so far by these initiatives and, more importantly, what still needs to be done.

### CLEF 2010: The Labs and The Workshops

The laboratories continue and expand on the traditional CLEF track philosophy. Two different forms of labs are offered: benchmarking activities proposing evaluation tasks and comparative performance analyses, and workshop-style labs that explore issues of information access evaluation and related fields.

There were five benchmarking activities:

- CLEF-IP '10, sponsored by the Information Retrieval Facility (IRF) in Vienna investigated patent retrieval in a multilingual context.
- ImageCLEF 2010 offered several tasks for both context and content-based retrieval of medical images from the journal of Radiology and Radiographics and photographic images from Flickr and Wikipedia.
- PAN, sponsored by Yahoo Research, investigated the detection of plagiarism and of Wikipedia vandalism.
- QA@CLEF 2010 used the Corpus of the European parliament and offered monolingual question answering tasks for English, French, German, Italian, Portuguese, Spanish and Romanian.
- WePS (Web People Search) focused on person name ambiguity and person

attribute extraction and online reputation management

And two exploratory workshops:

- CriES addressed the problem of multi-lingual expert search in social media environments.
- LogCLEF studied search behaviour in a multilingual context via query analysis and classification.

CLEF 2010 was hosted by the Department of Information Engineering, University of Padua, Italy and was partially supported by the PROMISE project. Approximately 140 researchers participated in the event; with most of them staying for the full four days. All the presentations given at the conference and in the workshops can be found on the CLEF 2010 website

### CLEF 2011

CLEF 2011 will be hosted by the University of Amsterdam in September 2011. A Call for Lab proposals will be issued at the beginning of October 2010 and an initial Call for Papers will be released in December 2010.

### Links:

CLEF 2010: <http://www.clef2010.org/>  
PROMISE: <http://www.promise-noe.eu/>

### Please contact:

Nicola Ferro  
University of Padua, Italy  
E-mail: [ferro@dei.unipd.it](mailto:ferro@dei.unipd.it)



## HCI International 2011 14th International Conference on Human-Computer Interaction

Hilton Orlando Bonnet Creek, Orlando, Florida, USA,  
9-14 July 2011

*Areas of Interest:* Ergonomics and Health Aspects of Work with Computers; Human Interface and the Management of Information; Human-Computer Interaction; Engineering Psychology and Cognitive Ergonomics; Universal Access in Human-Computer Interaction; Virtual and Mixed Reality; Internationalization, Design and Global Development; Online Communities and Social Computing; Augmented Cognition; Digital Human Modeling; Human Centered Design; Design, User Experience, and Usability. For more information about the topics listed under each thematic area, please visit the Conference website.

*Awards:* A plaque and a certificate will be given to the best paper of each of the eleven Affiliated Conferences / Thematic Areas, amongst which, one will be selected to receive the golden award as the Best HCI International 2011 Conference paper. Moreover, the Best Poster extended abstract will also receive a plaque and a certificate.

*Keynote address:* Ben Shneiderman (Professor in the Dept of Computer Science, founding director of the HCI Lab and Member of the Institute for Advanced Computer Studies at the University of Maryland) will give the keynote address "Technology-Mediated Social Participation: The Next 25 Years of HCI Challenges".

*Proceedings:* The Conference Proceedings will be published by Springer in a multivolume set in the LNCS and LNAI series, and will be available on-line through the LNCS Digital Library.

### Summary of Submission Requirements & Deadlines

- Abstract length, deadline for abstract receipt:
- Papers (800 words): Friday, 15 October 2010
- Posters (300 words): Friday, 11 February 2011
- Tutorials (300 words): Friday, 15 October 2010

### More information:

<http://www.hcii2011.org>

### Please contact:

Constantine Stephanidis  
FORTH-ICS, Greece  
E-mail: [cs@ics.forth.gr](mailto:cs@ics.forth.gr)



Boulder, Colorado, USA, 20-24 June 2011

CompArch is a federated conference series bringing together researchers and practitioners from Component-Based Software Engineering and Software Architecture. It features a joint program on component-based systems, software architectures, and their quality characteristics.

The federated events within **CompArch 2011** are: the Symposium on Component-Based Software Engineering (**CBSE**), the Conference on the Quality of Software Architectures (**QoSA**), the Symposium on Architecting Critical Systems (**ISARCS**), the Workshop on Component-Oriented Programming (**WCOP**), and the Working Conference on Software Architecture (**WICSA**)

### More information:

<http://comparch2011.archspot.com/>



## Multilingual Web Project Workshop

Madrid, 26-27 October 2010

The MultilingualWeb Project, funded by the European Commission and coordinated by the W3C, is looking at best practices and standards related to all aspects of creating, localizing and deploying the multilingual Web. The project will raise visibility of what's available and identify gaps via a series of four events, over two years. This first Workshop is free and open to the public.

Speakers will represent a wide range of organizations and interests, including: BBC, DFKI, European Commission, Facebook, Google, Loquendo, LRC, Microsoft, Mozilla, Opera, SAP, W3C, WHO, and the World Wide Web Foundation. Session titles include: Developers, Creators, Localizers, Machines, and Users. The Workshop should provide useful cross-domain networking opportunities.

### More information:

<http://www.w3.org/International/multilingualweb/madrid/program>



Call for Participation

## 2nd International ICST Conference on Cloud Computing

Barcelona, Spain, 25-28 October 2010

Cloud computing is an emerging computing paradigm envisioned to change all IT landscape facets including technology, business, services, and human resources. It is a consumer/delivery model that offers IT capabilities as services, billed based on usage. Many such cloud services can be envisioned, but the main ones are IaaS (Infrastructure-as-a-Service), PaaS (Platform-as-a-Service), and SaaS (Software-as-a-Service).

The underlying cloud architecture includes a pool of virtualized compute, storage, and networking resources that can be aggregated and launched as platforms to run workloads and satisfy their Service-Level Agreement (SLA). Cloud architectures also include provisions to best guarantee service delivery for clients and at the same time optimize efficiency of resources of providers. Examples of provisions include, but are not limited to, elasticity through scaling resources up/down to track workload behavior, extensive monitoring, failure mitigation, and energy optimization. The two main technologies enabling clouds are: (i) virtualization, the foundation of clouds; and (ii) manageability (autonomics), the command and control of clouds. CloudComp 2010 is intended to bring together researchers, developers, and industry professionals to discuss clouds, cloud computing, and related ecosystem support.

The technical program of the conference will include the following areas:

- Cloud architectures and provisions to optimize providers' environments while guaranteeing clients' SLAs
- Programming models, applications and middleware suitable for dynamic cloud environments
- End-to-end techniques for autonomic management of cloud resources including monitoring, asset management, process automation and others
- New cloud delivery models, models' optimizations and associated architectural changes
- New cloud economic and billing models
- Cloud security, privacy and compliance challenges
- Toolkits, frameworks and processes to enable clouds and allow seamless transitions from traditional IT environments to clouds
- Experiences with existing cloud infrastructure, services and uses
- Novel human interfaces and browsers for accessing clouds
- Interaction of mobile computing, mCommerce and clouds.

**More information:**

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



Call for Participation

## FET11 - The European Future Technologies Conference

Budapest, 4-6 May 2011

Following on from an exciting FET event in Prague last year, the DG Information Society and Media of the European Commission has announced that the next FET conference and exhibition will take place in Budapest on 4-6 May 2011, during the Hungarian presidency of the European Union.

FET11 will display "Science beyond fiction" with an impressive list of keynote speakers and futuristic exhibition booths, covering a broad range of scientific fields and advanced technologies enabled by ICT. Calls for sessions, exhibition and posters will soon be released. In the meantime, have a look at [http://ec.europa.eu/information\\_society/events/fet/2009/](http://ec.europa.eu/information_society/events/fet/2009/) to get a flavour of what FET11 will bring to you!

The 2011 edition of this new European forum dedicated to frontier research in future and emerging information technologies is organised in collaboration with a coordination action led by ERCIM and including SZTAKI, the Hungarian member of ERCIM.

**More information available shortly at:**

<http://www.fet11.eu/>

## Sylvain Lefebvre Winner of the Eurographics Award 2010



*Sylvain Lefebvre.*

Sylvain Lefebvre, researcher at INRIA, has won the Eurographics Award in the category 'Young Researcher'. This award comes in recognition of his research work on texture synthesis. He views this award as acknowledgement of the collective work carried out within the teams he has worked with at INRIA on algorithms and methods to facilitate the creation and display of virtual environments. Sylvain Lefebvre has focused his research on textures and

the automated creation of graphic content with the aim of improving the quality of interactive environments. "Texture" is the method used to give an appearance to the surface of virtual objects. Textures complement geometric modelling, which defines the shape of virtual objects. Sylvain Lefebvre and his colleagues have attempted to answer questions such as how colour points can be attached to a surface to imitate a material such as wood, and how colours can be composed automatically by an algorithm to represent stone or paper, using the 'example-based synthesis method'.

### More information:

<http://www.inria.fr/actualites/2010/eurographics2010.en.html>

## W3C UK and Ireland Office Moves to Nominet

After 13 years of successful work at STFC Rutherford Appleton Laboratory, the W3C UK and Ireland Office has a new home at Nominet. Nominet runs the one of the world's largest Internet registries, the registry for .uk domain names, with over eight million domain names. Phil Kingsland, Director of Marketing and Communications at Nominet, will be the new Office manager. He said, "We believe the work W3C does promoting Web accessibility standards, and developing other standards that help Web users to trust in the reputation of the Internet is well aligned with Nominet's public purpose remit and vision, which is to be a leading force in making the Internet a trusted space, which everyone can be part of and has a positive impact on people's lives." The Office plans a ceremonial launch later this year.

W3C would like to thank STFC Rutherford Appleton Laboratory and the W3C UK and Ireland Office staff, led by Michel D Wilson and his predecessors Stuart Robinson and Bob Hopgood, for their contributions to W3C and the Web. Learn more about the W3C Offices, regional W3C representatives that help promote the W3C mission.

### More information:

<http://www.w3.org/Consortium/Offices/>

## Peter Bosman wins Best Paper Award at Genetic and Evolutionary Computation Conference 2010

Peter Bosman, researcher at CWI has won a Best Paper Award during the Genetic and Evolutionary Computation Conference 2010 in Portland, Oregon (USA). This is one of the most outstanding conferences on Evolutionary Computation. Bosman received the prize for his publication 'The Anticipated Mean Shift and Cluster Registration in Mixture-based EDAs for Multi-Objective Optimization'. He

*Peter Bosman (left) receives the best paper award.*



won the prize in the category of Estimation of Distribution Algorithms. Estimation of Distribution Algorithms (EDAs) are advanced genetic algorithms that are mainly used for the solution of general optimization problems. Bosman's winning article studies the characteristics of probability distributions and their influence on the optimization by EDAs. EDAs are broadly used, especially when it's difficult or impossible to use other techniques. Bosman focused on multi-objective optimization, where the goal is to optimize multiple, often conflicting objectives simultaneously, such as, for example, the costs and quality of a product. Bosman's research is part of the CWI research group Computational Intelligence and Multi-Agent Games (SEN4). EDAs were recently used during the study of adaptive bed planning in hospitals. They will likely be deployed in other projects, for example in research for revenue management and energy systems.

### More information:

<http://www.sigevo.org/gecco-2010/>



## Mobilize your Apps!

W3C created cards that summarize the Mobile Web Application Best Practices document. These guidelines aid the development of rich and dynamic mobile Web applications. This work is supported by the MobiWebApp FP7 EU project ([mobiwebapp.eu](http://mobiwebapp.eu)).

### More information:

<http://www.w3.org/2010/09/MWABP/>



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



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



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



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



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



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



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



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



Centrum Wiskunde & Informatica

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



Universitas Varsoviensis  
Universitas Wroclaviensis

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/>



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



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



Fonds National de la  
Recherche Luxembourg

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



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



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

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



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



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



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



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



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



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



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

## Order Form

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

Send, fax or email this form to:  
**ERCIM NEWS**  
**2004 route des Lucioles**  
**BP 93**  
**F-06902 Sophia Antipolis Cedex**  
**Fax: +33 4 9238 5011**  
**E-mail: [office@ercim.eu](mailto:office@ercim.eu)**

I wish to subscribe to the

**printed edition**

**online edition (email required)**

Name: .....

Organisation/Company: .....

Address: .....

Postal Code: .....

City: .....

Country: .....

E-mail: .....

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