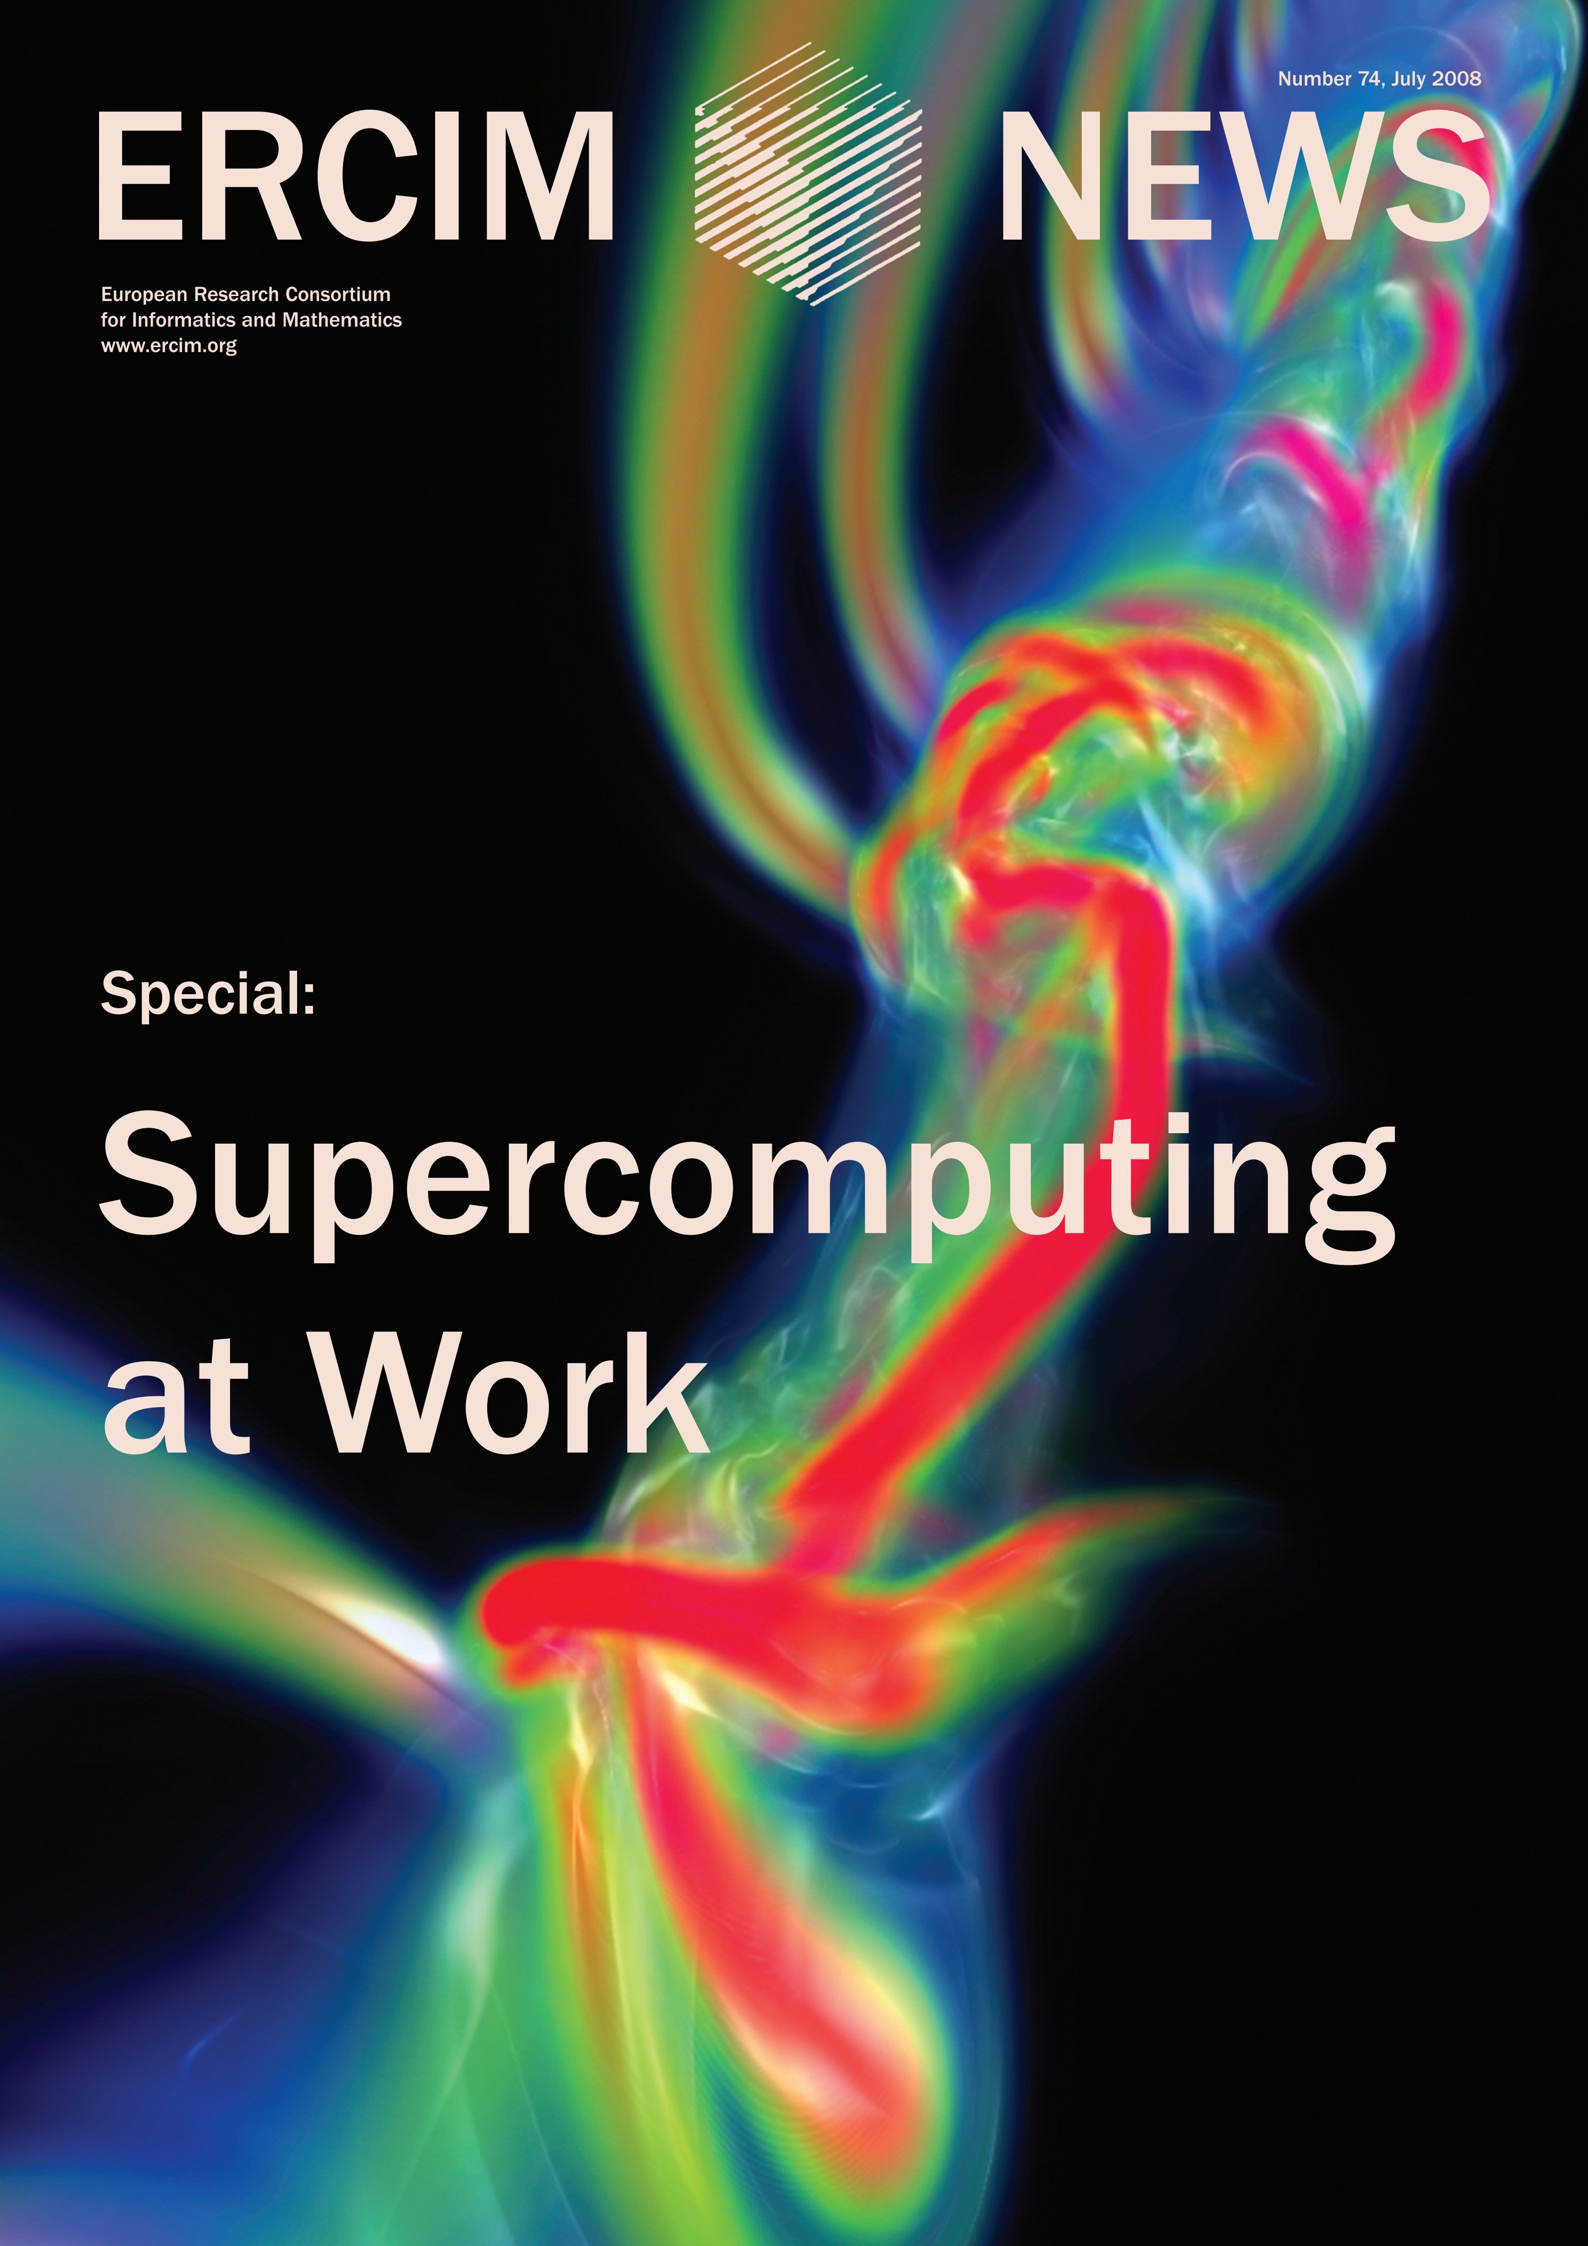

Number 74, July 2008

ERCIM NEWS

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

Special:

Supercomputing at Work

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

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

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

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

ERCIM News online edition
The online edition is published at http://ercim-news.ercim.org/

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

Next issue:
October 2008, Special theme: Safety-Critical Software

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

Credits:
Philippe Chatelain, Michael Bergdorf, Diego Rossinelli, Petros Koumoutsakos, ETH Zurich, Switzerland
Alessandro Curioni, Wanda Andreoni, IBM Zurich Research Laboratory, Switzerland.

Acknowledgments:
IBM T.J. Watson Research Center
Swiss Supercomputing Center.

Knocking at Petaflop's Door

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

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

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

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

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

Fortunately there is good reason to be optimistic that the simulation sciences – often denoted as the third column of knowledge creation – will also flourish over here. Support for supercomputing infrastructures has been established in the 7th Framework Programme, and following the late 2007 recommendations of the European Strategy Forum on Research Infrastructures, the European Commission has put the creation of a European supercomputing infrastructure on its agenda. These top-level systems will provide levels of performance comparable to other top installations worldwide for the benefit of Europe's scientific and technological development.

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

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

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

Thomas Lippert

Next issue: October 2008
Next special theme: Safety-Critical Software

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

However this is not simply another large-scale simulation problem, since the commercially sensitive nature of the problems and stringent IP protection demands fine-grained security on access to, and usage of, licensed software, providing protection for the intellectual property associated with circuit and device designs, data and simulations belonging to industrial partners and key stakeholders. To this end, the project has developed infrastructure that provides security throughout. This includes exploitation of Kerberos for secure global file-based access through the Andrew File System; authorization technology such as PERMIS for definition and enforcement of access policies using centralized attribute authorities such as the Virtual Organization Membership Service (VOMS); and simple user-oriented access to the project portal through the Internet2 Shibboleth technology using the UK Access Management Federation. This portal contains a variety of services for atomistic device simulation, compact model generation and circuit simulation, as well as mechanisms for managing the large and heterogeneous data and metadata associated with these simulations.

The project is now well advanced and the initial prototype infrastructure is ramping up for large-scale scientific usage. Up until now, due to the computational complexity of 3D device simulation, studies of variability have tended to be based on small ensembles of typically up to 200 devices. We have simulated ensembles in excess of 100,000 3D devices for 35nm and 13nm gate-length devices. Ensembles of this magnitude are shedding new light on the effects of atomic structure variation on the behaviour of devices, especially extreme limits of device variability. Furthermore, based on these simulations, we have been able to examine the effect of device variability at a simple circuit level and have simulated over 1 million CMOS inverters using random configurations of devices. Figure 1 shows the potential and dopant position of a statistically rare device. Figure 2 shows the threshold voltage variation as a function of the number of dopants.

The second phase of this project will use the methodologies developed to study larger and more advanced circuits and systems and to further explore the impact of atomistic variability of transistors on the design process. More information on the nanoCMOS project is available on our Web site.

Links:
UK Engineering and Physical Sciences Research Council:
http://www.epsrc.ac.uk
The nanoCMOS project:
http://www.nanocmos.ac.uk
UK e-Science National Grid Service:
http://www.ngs.ac.uk
ScotGrid: http://www.scotgrid.ac.uk

Please contact:
Asen Asenov
(science-related questions)
University of Glasgow, UK
E-mail: a.asenov@elec.gla.ac.uk

Richard Sinnott
(e-Infrastructure-related questions)
University of Glasgow, UK
E-mail: r.sinnott@nesc.gla.ac.uk

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

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

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

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

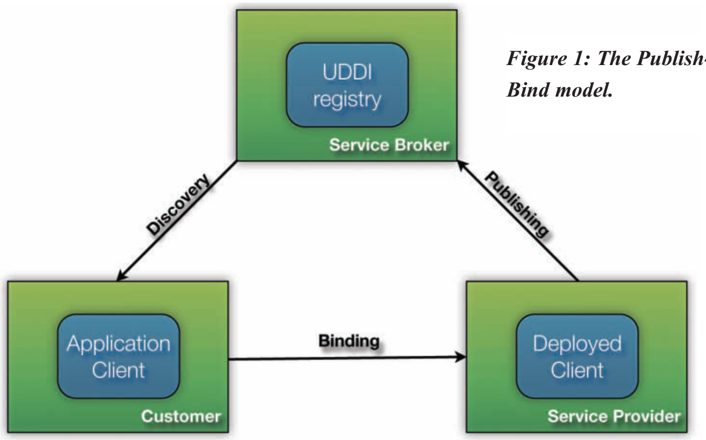

Figure 1: The Publish-Discover-Bind model.

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

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

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

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

Our Contribution

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

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

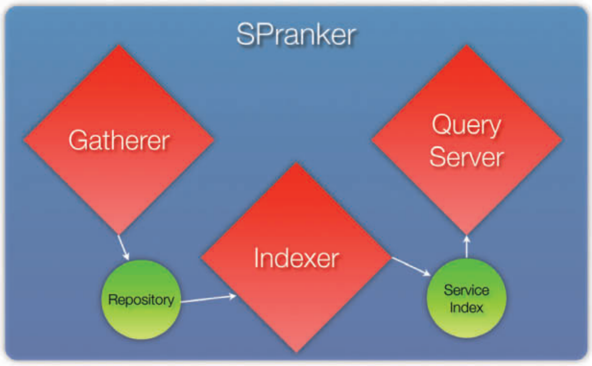

Figure 2: Architecture of SPRanker.

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

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

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

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

Assume a Query-SLA and a set of SLA-vectors representing an SLA successfully issued at time T, by a provider, on a particular service. A similarity function that takes into account the provider, service name and SLA issue time is defined. We define sim=0 if either the provider or the service name differs. Otherwise the similarity is defined as the sum of the common terms shared by Query-SLA and an SLA-vector weighted by the time elapsed since the SLA was issued. Presenting a Query-SLA, SPRanker seeks a list of providers offering a particular service ordered by similarity with the query.

The architecture of SPRanker is composed of three modules: gatherer, indexer and query server. The gatherer collects data from (positively) expired SLAs. We only consider positively expired SLAs because we want to discriminate between good and bad service provisioning, and because we want to avoid satisfied customers incurring false bad judgements from malicious partners (clients or customers) willing to lower a provider's score. The gatherer can act in two different modes – push-based and pull-based.

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

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

Link:
http://www.s-cube-network.eu/

Please contact:
Franco Maria Nardini
ISTI-CNR, Italy
Tel: +39 0050 3153055
E-mail: francomaria.nardini@isti.cnr.it