

A virtual stone soup: letter from the editors

Luís Veiga · Fabio Kon · Gordon Blair

Published online: 24 August 2012
© The Brazilian Computer Society 2012

This issue of the Journal of Internet Services and Applications features six papers covering the fields of Cloud/Grid/Autonomic Computing, Recommendation Systems, Web Services, and Security. It is released when the greatest media event on Earth has just taken place: the Olympic Games of London 2012. The experience of The Games has been made more pervasive, in-depth, and personalized as never before with fans being able to follow their favorite sports on a multitude of hand-held devices, either live or asynchronously. An unprecedented computing infrastructure has been set up to produce and deliver, via Internet, content (schedules, news, records and result history, athlete profiles) and live feeds of most events, even those that are not to be broadcasted by any TV station. Cloud computing and cloud-enabled technologies were used for capturing, storing, editing, and enhancing content, such as projected flags on tracks and swimming pools, lines highlighting record marks in pools and sandboxes for track and field events.

Encoding and streaming live and on-demand has required massive computation, storage, and network deployment. But, besides the sheer scale of equipment and deployment, there are also challenges and opportunities related to data handling and processing, as well as application integration.

We thus see that this intensive use of computer-related technologies and, in particular, of Internet-related

technologies in the coverage of this Olympiad puts into practice all the science produced by the Internet research community to improve the experience of billions of people in their enjoyment with the games. Current research challenges and opportunities will certainly lead to more technology that will help to improve even more the coverage of future olympic games.

For this issue and the next, the editors-in-chief counted with the proactive help of Prof. Luís Veiga from INESC-ID Lisboa, Portugal who helped to collect some of the best works presented in the Middleware'2010 workshops; they are published in these two issues of JISA as extended and enhanced papers. For that, all workshop chairs were asked, according to the number of published papers in the workshop, to select one or two papers for possible journal publication. The authors of selected papers were invited to submit a thoroughly revised, updated, and extended version of their work to JISA. After an additional round of reviews, three papers were eventually selected for publication: one in this number and two in the next one.

The first of the selected papers, “Tuning adaptive computations for the performance improvement of applications in JEE server” by Ying Zhang, Gang Huang, Xuanzhe Liu, and Hong Mei, selected from the 9th International Workshop on Adaptive and Reflective Middleware (ARM'2010), handles the issues brought by adaptation itself that, while increasing flexibility and application efficiency, also takes its toll on resources when it is activated. The authors help on finding out when adaptation is paying off or is simply adding load to the system. The help comes in the shape of an adaptation tuning model designed to improve the performance of adaptive applications, implemented for JEE servers, by upgrading or downgrading autonomic levels, according to available resources. The tuner is used to optimize Pkuas and JOnAS, two autonomic JEE servers.

Luís Veiga
INESC-ID Lisboa, Instituto Superior Técnico,
UTL, Rua Alves Redol 9, 1000-029 Lisbon, Portugal

Fabio Kon (✉)
University of São Paulo, Rua do Matão,
1010, São Paulo, SP 05508-090, Brazil
e-mail: fabio.kon@ime.usp.br

Gordon Blair
Lancaster University, South Drive, Lancaster LA1 4YW, UK

Extensive benchmark evaluation is carried out with the ECperf and RUBiS benchmarks.

Among the regular papers submitted to JISA, we start with a paper selected by our guest editor, Dr. Dilma Da Silva, “Towards an opportunistic grid scheduling infrastructure based on tuple spaces” by Fábio Favarim, Joni da Silva Fraga, and Lau Cheuk Lung. In this paper, the authors propose a new scheduling infrastructure for grids where resources select the tasks they execute, instead of the traditional approach in which schedulers find resources for the tasks. With this new approach, the system can make scheduling decisions with more up-to-date and accurate information about resource availability.

Cloud computing has become almost prevalent with many players providing utility-like computing resources with economies of scale and close-to-unlimited elasticity. Yet, after the fairy-tale, users and customers find out not all clouds are alike, especially regarding businesses migrating their applications to the cloud. Basem Suleiman, Sherif Sakr, Ross Jeffery, and Anna Liu, in “On understanding the economics and elasticity challenges of deploying business application on public cloud infrastructure”, bring order to the chaos by analyzing and discussing, with motivating scenarios, the economics and elasticity challenges that business applications face when being deployed to operate in a public cloud infrastructure. The authors model and quantify elasticity in settings with different workload patterns, policies, and SLAs.

Context-awareness, a classic middleware topic, is revisited by Kam Fung Yeung, Yanyan Yang, and David Ndzi in “A proactive personalised mobile recommendation system using analytic hierarchy process and Bayesian network”, this time in the frame of recommendation systems for mobile users. Many of today’s Internet users are almost always online but the immense amounts of web content make discerning the wheat from the chaff a daunting task. Automatic recommendation systems help by suggesting content that will interest the user with a high probability, based on her previous browsing and profile. The authors leverage Bayesian

networks processing activities, location, usage patterns to establish context information. Then, recommendations can be ranked with multi-criteria, with all the recommendation system exposed to programmers via an API.

In “A robustness testing approach for SOAP Web services”, Nuno Laranjeiro, Marco Vieira, and Henrique Madeira address the very important, but understudied, topic of testing of distributed systems. Testing practices are nowadays highly used in the industry for the development of complex software systems. However, most of the testing research as well as technologies and tools are targeted at centralized systems. This paper addresses the part of this problem by focusing on the robustness testing of Web services. The approach proposed by the authors is based on a set of robustness tests, including both malicious and non-malicious invalid call parameters, which are used to discover programming and design errors. The proposal is validated by two sets of experiments showing the use of Web services robustness testing from the consumer and the provider points of view. The experiments comprise the robustness testing of over 1,200 public Web service operations available in the Internet as well as 29 services implemented by the authors.

Last but not least, Wei-Chuen Yau, Raphael Phan, Swee-Huay Heng, and Bok-Min Goi present “Security models for delegated keyword searching within encrypted contents”. In the Internet, it is often needed to delegate operations across nodes without compromising the confidentiality of information. In this paper, the authors present new security models to capture keyword guessing attacks against public key encryption with keyword search. Their new models allow to capture attacks that were not resolved in previous schemes.

As inspiring and diverse as the Olympic Games, we expect that the papers of this issue will serve as a delicious and nourishing pot of stone soup¹ enjoyed by all.

Luís Veiga, Fabio Kon, Gordon Blair
August/2012

¹ http://en.wikipedia.org/wiki/Stone_soup.