

## The path to openness: letter from the editors

**Fabio Kon · Gordon Blair · Lisandro Granville ·  
Raouf Boutaba · Lizhe Wang · Jie Tao ·  
Achim Streit · Dieter Kranzlmüller · Luís Veiga**

Published online: 1 November 2012  
© The Brazilian Computer Society 2012

With this final JISA issue in 2012, we come to the end of the first phase of our journal. Since May 2010, when the first issue was published, JISA has encountered a growing acceptance in the academic and industrial research communities in topics related to the Internet. Indeed, the journal is now blossoming. The number of paper submissions have been increasing significantly and the impact of the papers published in the journal is growing.

It is our great pleasure to announce a second phase of JISA with the journal becoming an Open Access Journal, starting on January, 2013! We will be a part of Springer Open and

all JISA papers will be freely accessible to anyone across the globe at <http://www.jisajournal.com>. The copyright of the papers will remain with their authors. We are very happy with the new configuration of the journal as it democratizes access to the information, which we expect to be valuable for our community.

To end the current phase and start a new one, the current issue combines four different sets of papers in a bumper issue: on Network Virtualization guest-edited by Lisandro Granville and Raouf Boutaba, on Data-Intensive Computing guest-edited by Lizhe Wang, Jie Tao, Achim Streit, and Dieter Kranzlmüller, two papers selected from the Middleware'2010 best workshop papers, organized in collaboration with Luís Veiga, and, finally, three regular papers.

---

F. Kon (✉)  
University of São Paulo, São Paulo, Brazil  
e-mail: fabio.kon@ime.usp.br

G. Blair  
Lancaster University, Lancaster, UK

L. Granville  
Federal University of Rio Grande do Sul,  
Porto Alegre, Brazil

R. Boutaba  
University of Waterloo, Waterloo, Canada

L. Wang  
Center for Earth Observation, Chinese Academy of Sciences,  
Beijing, China

J. Tao · A. Streit  
Karlsruhe Institute of Technology,  
Karlsruhe, Germany

D. Kranzlmüller  
Ludwig Maximilians Universität München,  
Munich, Germany

L. Veiga  
Instituto Superior Técnico-UTL/INESC-ID, Lisboa, Portugal

### Regular papers

We start with the three regular papers. Tiago Almeida and Akebo Yamakami present “Occam’s Razor-based Spam Filter,” featuring a novel approach to spam filtering based on the minimum description length principle. Ioannis Papadakis, Michalis Stefanidakis, Sofia Stamou and Ioannis Andreou then present “Semantifying Queries over Large-scale Web Search Engines”, which proposes new techniques to enhance the quality of web search engines by taking semantics into consideration; their new service leverages information coming from various DBpedia datasets and provides an intuitive GUI via which searchers determine the semantic orientation of their queries before these are sent to the underlying search engine. The last regular paper, by Afonso Ferreira, is “Computing Multicast Trees in Dynamic Networks and the Complexity of Connected Components in Evolving Graphs”. In this paper, Ferreira uses evolving graphs to compute multicast trees with minimum overall transmission time for a class of wireless mobile dynamic networks.

### Middleware'2010 best workshop papers

In this issue, we finish publishing the extended and enhanced versions of the best workshop papers presented in the Middleware'2010 conference. We start with “A Reconfigurable Component Model with Semantic Type System for Dynamic WSN Applications” by Klaas Thoelen et al., which presents a novel component model for the development of wireless sensor network (WSN) applications operating in dynamic environments and under evolving application requirements. Then, in “Jano: Location-Privacy Enforcement in Mobile and Pervasive Environments through Declarative Policies”, José Simão et al. propose a middleware platform supporting both pull and push location requests while enforcing configurable, history-based, security and privacy policies.

### Network virtualization

In network virtualization, traditional computer networks become the substrate that hosts several virtual network infrastructures. Multiple virtual network nodes can run inside a single network physical device, such as a router or switch. As virtual nodes are isolated from one another, they can be created, moved, and destroyed across the physical substrate without affecting other nodes and the regular network operation. As such, network virtualization enables the disruptive experimentation of new network architectures without requiring dedicated, expensive physical devices. New network protocols can be deployed and tested in large-scale virtual network infrastructures. Diverse concurrent network solutions, such as content delivery network (CDN), information centric networking (ICN), and dynamic circuit network (DCN), devoted to different purposes, can coexist on top of the shared physical substrate, thus creating network virtualization environments (NVEs). The opportunities created with network virtualization lead network virtualization to be quite often referred to as the technology that can effectively tackle the so-called Internet ossification.

Three JISA papers will address important network virtualization questions, one in this current issue and two others in the open access format, early next year. The first paper of the theme, “A Distributed Controller for a Virtualized Router” by Mellah, Cherkaoui, and Lemieux, introduces a solution for the management of network routers with virtualization support employing a distributed controller.

### Data-intensive computing

Data-intensive computing refers to a type of computation, incorporating systems and applications, which uses parallel and distributed computing approaches to process large volumes of data, currently in the order of terabytes or petabytes. Data-intensive computing has become the fourth paradigm for scientific discovery, proposed by Turing laureate Dr. Jim Gray. Scientific advances in all fields of human knowledge are now heavily dependent on the aid of data-intensive computing. Recently, data-intensive computing has attracted significant research interest and great progress has been made in various aspects, such as hardware architectures, runtime systems and OS level support, programming models and middleware, storage systems, and applications.

In this issue, a selection of papers has been included to represent recent advances in various aspects of data-intensive computing. It is a key research issue to manage massive datasets for data-intensive computing applications. Wang et al. contribute to this area by developing a record-based block distribution (RBBD) framework and a weighted set cover scheduling for data-intensive analytics in MapReduce in the HDFS environments. Programming with massive data in a data-intensive computing environment is another challenging issue; in this regard, Costa et al. present GiGi-MR, a framework that allows non-expert users to run CPU-intensive jobs on top of volunteer resources over the Internet. Data-intensive applications bring real use cases to explain how the data-intensive computing paradigm works; Deng et al. present one such data-intensive application: a massively parallel Morlet continuous wavelet transform (MCWT) approach based on general-purpose computing on graphics processing units (GPGPU) to process massive electroencephalogram (EEG) data in a real-time manner. Early next year, a new set of papers on data-intensive computing will be published in the open access format.

We hope you enjoy this last issue of JISA in its current form, as it is filled with interesting new ideas from a variety of topics. Next year, with its new Open Access format at <http://www.jisajournal.com>, JISA will continue to display high-quality papers but with an increased potential for impact and knowledge dissemination.