



Contents lists available at SciVerse ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Editorial

Special section: Recent advances in utility and cloud computing

1. Introduction

Welcome to the special issue of Future Generation Computer System (FGCS) Journal in Utility and Cloud Computing. This special issue compiles a number of excellent technical contributions that significantly advance the state-of-the-art in utility and cloud computing comprising selected and extended contributions from the UCC 2010, the 2010 International Conference on Utility and Cloud Computing, held in Chennai, India, in December 2010. Papers selected for UCC 2010 were peer-reviewed by 2–3 reviewers. Authors have improved their papers based on conference reviews and presented them at UCC 2010 in India. Based on the feedback authors received at the conference, they have further extended their papers and submitted them to FGCS. All papers submitted to FGCS are further reviewed to meet its high standard expectation. Authors were requested to revise their papers to address comments from FGCS reviewers and to resubmit them. These revised papers were finally accepted by the Editor-in-Chief of FGCS for publication in this special issue.

Selected papers reflect the current trends and challenges in establishing reliable, self-adaptable, and economy-driven Clouds and Utility computing infrastructure, platforms, and applications. It is important to note that even the most sophisticated Cloud infrastructure can fail to reach the critical consumer mass, if the market mechanisms do not support the needs of the market participants. Thus, the first trend considers establishment of *Cloud markets and generation of Service Level Agreements between Cloud market participants*. Mega scale and heterogeneous Cloud infrastructure needs reliable *Cloud resource management* considering large data transfers, advance resource reservation and sophisticated autonomic management of resources representing the second trend we report in this special issue. On the one hand, Cloud computing has a huge potential to revolutionize IT infrastructure by providing novel business models. On the other hand, due to virtualization overheads and mega-scale infrastructure there is potential for energy waste. The third trend we focus on in this special issue is *(energy) efficient usage of resources & solutions for green Clouds*. The fourth topic we cover in this special issue deals with *Cloud programming models* and application of traditional Cloud programming models (e.g., MapReduce) to various application areas, e.g., scientific computing. We will discuss how papers, selected for this special issue, address these four key topics.

2. Emerging topics

(1) *Cloud markets/SLA management*. Maurer *et al.* presents a novel approach for the management of liquidity in Cloud markets

by introducing the SLA mappings where users can express their difference to the public SLA templates used by market instruments. By applying the SLA mapping approach the user's utility to find an appropriate match in the market is increased, the balance between utility gain and costs for performing mappings are in a reasonable balance—thus, increasing the market liquidity. Ralf Vigne *et al.* discusses a novel approach for the management of a hybrid service and process repository in order to enable the dynamic selection and consumption of functionalities in Clouds increasing the chance that small providers can sell their Cloud products in the market. The approach is evaluated with a set of generic use cases according to customers' goal functions. Johan Tordsson *et al.* discusses goals, objectives, and challenges of the OPTIMIS project designing a holistic approach to Cloud service provisioning focusing on the set of tools for the optimization of the whole service life cycle, including service construction, deployment, and operation, on a basis of aspects such as trust, risk, eco-efficiency and cost. Yohei Matsushashi *et al.* propose a transparent VPN failure recovery scheme that can hide VPN failures from users and operating systems (OSs). Thereby, whenever a VPN failure occurs, a client virtual machine monitor (VMM) automatically reconnects to an available VPN gateway, which is geographically distributed and connected via leased lines in Clouds.

(2) *Cloud resource management*. Antonio Celesti *et al.* address the problems of transferring huge amounts of data among federated Clouds. They propose an approach based on satellite communications. The authors apply the proposed strategy in a real scenario, specifying the involvement of Cloud providers, Cloud users, satellite companies and end-user clients. Amit Nathani *et al.* present Haizea, a Cloud infrastructure manager that supports four kinds of resource allocation policies: immediate, best-effort, advanced reservation and deadline sensitive, thus facilitating a better way to support deadline sensitive leases while minimizing the total number of rejected leases. Hamoun Ghanbari *et al.* discusses the allocation of resources in a private Cloud such that the cost of the providers is minimized through the maximization of resource sharing while attempting to meet all client application requirements as specified in the SLAs. The presented allocation algorithm works in two modes: one that relates the application level quality of service to the given set of resources and one that maps a given service level and resource consumption to profit metrics. Pallickara *et al.* discuss a novel approach for discovering the correct dataset in an efficient way for computational science applications. Compared to text-based web documents, many of the large scientific datasets contain

binary or numerically encoded data that is hard to discover through the popular search engines. In atmospheric sciences, there has been a significant growth in public data hosting. The presented Atmospheric Data Discovery System (ADDS) provides an efficient data discovery environment for the observational datasets in the atmospheric sciences. To support complex querying capabilities fine grained meta data are automatically extracted and indexed.

- (3) *Green clouds/ resource efficiency.* Ericson *et al.* present the framework Granules, which provides Cloud developers use of the runtime's native language for computations. Granules handles a number of different programming languages by supporting multiple methods of communication such as named pipes, unnamed pipes, and sockets. This choice of underlying communication mechanisms allows limited resources, such as sockets, to remain available for use by the runtime. Cloud runtimes are an effective method of distributing computations, but can force developers to use the runtime's native language for all computations. Jérôme Gallard *et al.* discusses an architecture for the next generation system management tools by advocating a system management framework for the automatic adaptation of both, hardware and software resources to the applications' needs. Thereby the authors apply the concept of Virtual Platform (VP) and a Virtual System Environment (VSE). The formalisms for the description of VPs and VSEs are based on the Goldberg theory and novel concepts based on system virtualization (identity, partitioning, aggregation) and emulation (simple, abstraction). Rodrigo Callheiros *et al.* presents Aneka's deadline-driven provisioning mechanism responsible for supporting QoS-aware execution of scientific applications in hybrid clouds composed of resources obtained from a variety of sources. Thus, Aneka is capable of provisioning resources obtained from a variety of sources including private and public clouds, clusters, grids, and desktops grids.
- (4) *Cloud programming models.* Suraj Pandey *et al.*, design a novel system for a scalable and economic real-time health monitoring and analysis for people who require frequent monitoring of their health. The system design is focused on aspects of an autonomous Cloud environment that collects peoples health data and disseminates them to a Cloud-based information repository and facilitates analysis on the data using software services hosted in the Cloud. The system is evaluated by means of an experimental testbed on a specific use-case, namely, the collection of electrocardiogram (ECG) data obtained at real-time from volunteers to perform basic ECG beat analysis. Sadeka Islam *et al.* discusses prediction-based resource measurement and provisioning strategies using neural networks and linear regression to satisfy upcoming resource demands. Thereby the authors address efficient initialization of new virtual instances in clouds which may introduce several minutes' delay in the hardware resource allocation. Experimental results demonstrate that the proposed technique offers more adaptive resource management for applications hosted in cloud environment, an important mechanism to achieve on-demand resource allocation in the cloud. Eugen Feller *et al.* present the outcome of the EU-funded XtreamOS project that implements an open-source grid operating system based on Linux. In order to provide fault tolerance and migration for grid applications, it integrates a distributed grid-checkpointing service called XtreamGCP. The service is designed to support various checkpointing protocols and different checkpoint packages (e.g. BLCR, LinuxSSI, OpenVZ, etc.) in a transparent manner through a uniform checkpoint interface. Arijit Mukherjee *et al.* investigate the distributed query processing, as for example, OGSA-DAI and OGSA-DQP, which provide common interface to heterogeneous databases, and a way of exploiting

distributed resources. The authors propose an approach for dynamically deploying query processing engines, analysis of services and databases within virtual machines, on an Internet scale, so as to reduce communication costs. Results of Internet-scale experiments are presented to demonstrate the performance benefits. Satish Srirama *et al.* present a study about establishing private clouds at the universities. With these clouds, students and researchers can efficiently use the already existing resources of university computer networks, in solving computationally intensive scientific, mathematical, and academic problems. However, the existing applications have to be suited to Cloud programming models, e.g. to the MapReduce paradigm. The authors present an approach for reducing iterative algorithms to the MapReduce model. Spyridon V. Gogouvtis *et al.* discusses a new architectural design and implementation of a service-oriented workflow management tool by considering enacting soft real-time application service components according to a workflow description language, synchronizing the application components, monitoring the execution and reacting to events within a distributed virtualized environment. The authors also demonstrate the operation of the implemented mechanism and evaluate its effectiveness using an application scenario with soft real-time interactivity characteristics, namely film post-production, under realistic settings.

Acknowledgments

It is important to note that UCC 2010 did not publish pre-conference proceedings. All high quality papers selected for the UCC after peer review are revised and presented at the conference. Authors have further extended their papers based on feedback they received at the conference and submitted to FGCS Journal. These extended papers went to through further review via the FGCS review management system and have undergone further revision prior to their acceptance for this special issue, which serves as a "Post Conference Proceedings of UCC 2010". We would like thank all authors for their cooperation in the multi-phase reviews and revisions of their work.

We thank members of UCC Program Committee and reviewers of FGCS for their time and expertise in selecting high quality papers and commenting on improving the quality of all selected papers. We thank Thamarai Selvi from Anna University for providing the venue to host UCC 2010 in India.

We thank Peter Slood (Editor-In-Chief) of the FGCS journal for supporting us in publishing the UCC 2010 Proceedings as a special issue with rigorous quality control. We thank staff members of Elsevier, especially Rebecca Capone and Fred Kop, for their cooperation in getting the special issue out in a timely manner.

We hope that you will enjoy reading this special issue!

Pointer to Recent Texts

- R. Buyya, J. Broberg, and A. Goscinski (eds), *Cloud Computing: Principles and Paradigms*, Wiley Press, New York, USA, February 2011.



Dr. Ivona Brandic is Assistant Professor at the Distributed Systems Group, Information Systems Institute, Vienna University of Technology (TU Wien). Prior to that, she was Assistant Professor at the Department of Scientific Computing, Vienna University. She received her PhD degree from Vienna University of Technology in 2007. From 2003 to 2007 she participated in the special research project AURORA (Advanced Models, Applications and Software Systems for High Performance Computing) and the European Union's GEMSS (Grid-Enabled Medical Simulation Services) project. She is involved in the European Union's SCube project and she is leading the Austrian national FoSII (Foundations of Self-governing ICT Infrastructures) project funded by the Vienna Science and Technology Fund (WWTF). She is a Management Committee member of the European Commission's COST Action on Energy Efficient Large Scale

Distributed Systems. From June–August 2008 she was visiting researcher at the University of Melbourne. Her interests comprise SLA and QoS management, Service-oriented architectures, autonomic computing, workflow management, and large-scale distributed systems (Cloud, Grid, Cluster, etc.).



Dr. Rajkumar Buyya is Professor of Computer Science and Software Engineering; and Director of the Cloud Computing and Distributed Systems (CLOUDS) Laboratory at the University of Melbourne, Australia. He is also serving as the founding CEO of Manjrasoft., a spin-off company of the University, commercializing its innovations in Cloud Computing. He has authored 350 publications and four text books. He has also edited several books including "Cloud Computing: Principles and Paradigms" recently published by Wiley Press, USA. He is one of the highly cited authors in computer science and software engineering

worldwide (h -index = 53, g -index = 115, 15500 + citations).

Software technologies for Grid and Cloud computing developed under Dr. Buyya's leadership have gained rapid acceptance and are in use at several academic institutions and commercial enterprises in 40 countries around the world. Dr. Buyya has led the establishment and development of key community activities, including serving as foundation Chair of the IEEE Technical Committee on Scalable Computing

and five IEEE/ACM conferences. These contributions and the international research leadership of Dr. Buyya are recognized through the award of the "2009 IEEE Medal for Excellence in Scalable Computing" from the IEEE Computer Society, USA. Manjrasoft's Aneka Cloud technology developed under his leadership has received the "2010 Asia Pacific Frost & Sullivan New Product Innovation Award".

Guest editors

Ivona Brandic

Vienna University of Technology, Austria

E-mail address: ivona@infosys.tuwien.ac.at.

Rajkumar Buyya*

The University of Melbourne and Manjrasoft, Australia

E-mail address: raj@csse.unimelb.edu.au.

Available online 21 June 2011

* Corresponding address: Department of Computer Science and Software Engineering, The University of Melbourne, VIC 3053 Carlton, Melbourne, Australia.