



Editorial

Preface of special issue on the economics of computing services

The increase in take up and maturity of Cloud computing systems over recent years has led to the need for a greater emphasis on business and economic perspectives of Cloud systems. As Cloud computing involves leasing, outsourcing and purchasing of compute and data storage capacity, involving a variety of resources, providers of Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS) need to understand the market, economics, and revenue models in these emerging areas. A closely related theme is the recent focus on economic sustainability of computing and storage infrastructure, especially in the emerging financial downturn observed in many parts of the world. The challenge here is the understanding of the economic interdependencies between different types of resources, the legal frameworks, and the market approaches available. Additionally, as Cloud computing often involves sharing of internal business processes with external providers (in a de-perimeterized operating environment), security concerns also become key. This understanding is required for trading computing resources, bundling of SaaS and IaaS services, orchestration of SaaS services, establishing of contracts (i.e., service level agreements) between providers as well as between providers and end-users, and achieving seamless interoperability between services of different providers. Only if this interplay between services, legal frameworks, and market approaches is analyzed and understood, will a sustainable Cloud market emerge that will build the IT foundation for other industries in our economy.

The goal of this special issue is to present papers that address the above-mentioned topics. The papers originate from the Workshop on the Economics of Grids, Clouds, Systems, and Services (GECON). The GECON workshop series is intended to enable researchers and practitioners from academia, industry, and national research laboratories to identify business and economics-related issues, suggest solutions and discuss lessons learned and experiences gained. The topics addressed range from extensions to existing technologies and architectures, reports on deployments of technologies, empirical and theoretical economic analysis, to the introduction of new theoretical concepts. The purpose of this workshop is to present original work and build a strong community in this increasingly important research area.

Contributions which received the highest assessment during the GECON2010 review process, with 3–5 reviews per paper, showing the largest potential for impact, have been extended and significantly enhanced for this special issue. The resulting set of papers featured in this special issue also went through an additional thorough review process. Seven papers have been selected, covering a wide range of the above-mentioned research areas.

The paper entitled “A Recommender Mechanism for Service Selection in Service-Oriented Environments” by Konstantinos

Tserpes, Fotis Aisopos, Dimosthenis Kyriazis and Theodora Varvarigou suggest the use of ratings provided by multiple customers, who used a particular service, for supporting service selection. The authors identify how collaborative filtering techniques can be used to relate user ratings, and thereby group users with similar types of ratings for services, so that reputations of providers can be established. This results in the design of a service recommender mechanism and a service recommender protocol. Their work is validated using simulations.

The contribution, entitled “Supporting CPU-based Guarantees in Cloud SLAs via Resource-level QoS Metrics” by Inigo Gori, Ferran Julia, J. Oriol Fito, Mario Macias and Jordi Guitart proposes a resource-level metric for specifying CPU performance guarantees. It allows Cloud resource providers to dynamically allocate their resources among running services depending on demand. The customer’s CPU usage is used in the metric definition, while “fake” SLA violations are avoided when a customer’s task does not use all its resources. The work has been evaluated in a virtualized provider. The results show that their solution results in fewer SLA violations than other CPU-related metrics while maintaining QoS.

Gheorghe Cosmin Silaghi, Liviu Dan Serban, and Cristian Marius Litan investigate “A Time-Constrained SLA Negotiation Strategy in Competitive Computational Grids” within a generic framework for strategically negotiating service level values under time constraints. Their contribution makes use of an agent-based system utilizing Bayesian learning. Their work shows that setting time constraints may actually lead towards an optimal allocation of resources and fair satisfaction of participants. It enables players (participants) to cope with the limited duration of a negotiation session by quickly learning the required parameters.

Ioan Petri et al., in their contribution “Service Level Agreement as a Complementary Currency in Peer-to-Peer Markets”, identify how SLAs can be used as a “virtual” currency to support resource exchange within a distributed system. The work focuses on building trusted collaborations between resources users and providers by utilizing a virtual currency as the mechanism for establishing provider reputation. The greater the reputation a provider establishes within a market (based on successful adherence to their Service Level Agreements), the greater will be the value of their currency—and correspondingly, the greater their power to influence the market. They use a PeerSim-based simulation to demonstrate profit/loss that can arise within a market of collaborating peers, exchanging SLAs.

The contribution by Juthasit Rohitratana et al., entitled “Impact of Pricing Schemes on a Market for Software-as-a-Service and Perpetual Software”, focuses on developing an agent-based simulation to support the pricing of software licenses. The simulation, which allows modeling the interactions between

software buyers and vendors in a software market, compares Software-as-a-Service and perpetual software licenses under different pricing schemes. Four dynamic pricing schemes are used by the vendors: derivative-follower pricing, demand-driven pricing, skimming pricing, and penetration pricing. The software buyer responds to the pricing by selecting the most appropriate software license scheme based on four criteria (i.e., finance, software capability, organization, and vendor) using the Analytic Hierarchy Process (AHP) decision support mechanism. The authors show which of the four pricing schemes should be followed in a particular context.

The software license theme is continued in the paper by Claudio Cacciari, Daniel Mallmann, Csilla Zsigri, Francesco D'Andria, Björn Hagemeyer, Angela Rimpl, Wolfgang Ziegler and Josep Martrat entitled "SLA-based Management of Software Licenses as Web Service Resources in Distributed Computing Infrastructures". It focuses on supporting software license management as web service resources in distributed service oriented environments, decoupling authorization of license usage from authorization of application execution. This way, licenses can become mobile objects, traveling to the environment that demands execution of an application. The core technology, security aspects, and an evaluation of a prototype for software license management are presented.

Rodrigo N. Calheiros et al. introduced an architecture of "A Coordinator for Scaling Elastic Applications Across Multiple Clouds". This coordinator, which is capable of negotiating and trading resources with other Clouds, distributes applications to other Clouds,

in case of a sudden increase in local demand. Without such a coordinator, a Cloud provider will not be able to deliver uniform Quality of Service (QoS) in its own Cloud, violating SLAs. The authors validate their approach through a prototype implementation.

In preparing this special issue we would like to thank Peter Sloot, the Editor-in-Chief of the journal Future Generation Computer Systems (FGCS) for his support. Furthermore, we would also like to thank Geetha Lakshmi and Fred Kop of Elsevier and Adam Belloum for helping with all cases that needed specific consideration.

Jörn Altmann*

Seoul National University

Technology Management, Economics, and Policy Program
Department of Industrial Engineering, College of Engineering
599 Gwanak-Ro, Gwanak-Gu, Seoul 151-744, South Korea
E-mail address: jorn.altmann@acm.org.

Omer Rana

Cardiff University, UK

Rajkumar Buyya

University of Melbourne, Australia

Available online 30 March 2012

* Corresponding editor.