

Preface

The proceedings of the 11th International Conference on the Economics of Grids, Clouds, Systems, and Services (GECON) were held during September 16–18, 2014 in Cardiff, UK. Cloud computing has now become mainstream, whereas in 2003, when this conference was initiated, the general ideas behind Cloud computing were being talked about within the Web Services and Grid computing communities, with a number of (active) researchers also investigating the associated economic models. Whereas Grid computing was initially very much driven by the computational science and engineering communities, with data/computational infrastructure hosted by national laboratories and (generally) research organizations, Cloud computing has, from its beginnings, been promoted by industry. The aim of the GECON conference has been to bring together this technical expertise (e.g., in resource allocation and quality of service management) with economics expertise (focusing on both micro- and macro-economic modeling and analysis) in order to create effective solutions in this space. Many of the models proposed in this event over the years have now materialized, and its widened scope is more relevant than ever considering the recent developments in our service economy with respect to (automated) trading, pricing, and management of services. Eleven years later, GECON remains true to its original objectives and continues to focus on the marriage of these two types of expertise. In line with this combined focus, it is useful to see two different methodological approaches being considered at GECON, one focusing on a technical perspective to problem solving, the other on economic modeling. Authors adopting a more technical focus often use a quantitative approach to validate their findings, whereas those from economic modeling often adopt a mixed methods (qualitative/quantitative) approach. Therefore, a key objective of this conference series is also to promote this multi-perspective, multidisciplinary approach to validate findings and to introduce authors to the use of these ideas in their own research. The review process ensures that both perspectives are considered equally important and relevant.

GECON 2014 took place in Cardiff (the capital of Wales), a city and county borough whose history spans at least 6,000 years. The city has played an important role in the service industry in the UK, with Cardiff being the main port for exports of coal from the surrounding region for many years. To support the service industry, the city grew rapidly in the nineteenth century. Today, Cardiff hosts the National Museum of Wales, the National Assembly of Wales, and Cardiff University where the conference took place.

This year we received a number of high-quality paper submissions. Each submission received at least three reviews by members of an international Program Committee. Our final program consisted of five sessions (two of which are Work-in-Progress sessions). The schedule for the conference this year was structured to encourage discussions and debates – with a round-table session focusing on “Sharing Economy”. We believe such discussion sessions are essential to enable more open and informed

dialog between presenters and the audience, and to enable the presenters to better position their work for future events and to get an improved understanding of the impact their work is likely to have on the research community. The presentation sessions set up were:

Session 1: Cloud Adoption

Session 2: Work in Progress on Market Dynamics

Session 3: Cost Optimization

Session 4: Work in Progress on Pricing, Contracts, and Service Selection

Session 5: Economic Aspects of Quality of Service

Session 1 started with a paper by Kaufman, Ma, and Yu entitled “A Metrics Suite for Firm-Level Cloud Computing Adoption Readiness” [1], which focused on gauging perceptions of managers and organizations when considering Cloud computing for outsourcing IT infrastructure and data. The authors interviewed a number of companies to better assess the Cloud adoption readiness of these organizations from both a strategy/management and technology/operations perspective. The subsequent paper by Pallas on “An Agency Perspective to Cloud Computing” used a particular economic modeling approach (agency theory) to model potential conflicts (in issues such as security, legal compliance, long-term availability) between users and providers of resources within a Cloud system [2]. Two example scenarios were used to motivate the proposed approach.

Session 2 was a Work-in-Progress session and included two contributions, both focusing on how Cloud provisioning could be supported through aggregating capacity across multiple providers (likely to be of benefit for consumers) compared to accepting and using the capability from a single large provider. Kim, Kang, and Altmann presented their work on “Goliath vs. a Federation of Davids: Survey of Economic Theories on Cloud Federation”. It described the benefits and costs of federation for “small” Clouds and how “small” Clouds can compete with larger Cloud providers (with the latter often having an advantage through economies of scale) [3]. This contribution also discussed whether federation of smaller Clouds is economically viable and under what conditions such a federation could compete with larger providers. The contribution by Vega, Meseguer, and Freitag entitled “Analysis of the Social Effort in Multiplex Participatory Networks” investigated how “community Clouds” could be established by pooling together resources from different individual providers (“citizens”) in the context of community networks (essentially, networks established over a small geographical area using resources contributed to by citizens) [4]. The authors investigated message exchanges and traffic on forums and mailing lists associated with the Guifi.net community network to understand how a small number of users act as potential “social bridges” between members of the community. The objective of the work was to understand the impact such members have on the community as a whole and the potential for such members to act as hosts of a potential community Cloud.

Session 3 focused on understanding technical metrics associated with Cloud provisioning and how these could be optimized. The contribution by Lučanin et al. entitled “Energy-Aware Cloud Management through Progressive SLA Specification” investigated how virtual machines could be migrated across multiple data centers to provide a particular energy footprint while still maintaining the quality of service targets [5].

The authors considered workload profiles from the Wikipedia and the Grid 5000 projects to validate their work, demonstrating that by offering virtual machines with differing availability and prices they could achieve (on average) 39 % energy saving per virtual machine. The paper “CloudTracker: Using Execution Provenance to Optimize the Cost of Cloud Use” by Douglas et al. discussed how provenance information obtained from executed jobs on Amazon EC2 instances could be used to predict potential cost of execution [6]. The authors also demonstrated how such a job submission system could be integrated with an existing computational science application. The final paper in this session by Me entitled “Migration to Governmental-Cloud Digital Forensics Community: Economics and Methodology” discussed how government agencies (especially crime and policing agencies) could use Cloud computing infrastructure to carry out forensic investigations [7]. The paper compared the cost of using an in-house computational infrastructure with an outsourced, Cloud-based infrastructure for supporting crime investigations. A key focus was on large-scale data analytics and post event, forensic analysis.

Session 4 consisted of three Work-in-Progress papers focusing on pricing and contracts, especially how parameters such as insurance costs and penalties in Service Level Agreements (SLAs) can be characterized. The contribution by O’Loughlin and Gillam entitled “Performance Evaluation for Cost-Efficient Public Infrastructure Cloud Use” investigated performance variation that can be observed when mapping Amazon instances on physical processors [8]. The authors compared the performance of the same Amazon instance using a number of different benchmarks, showing the variability in performance when the instance was requested over different times and within different data centers (using different physical processors to host the instances, e.g., use of Intel vs. AMD). This variability clearly shows the difficulty of developing SLAs when using public Cloud infrastructure. The contribution by Naldi entitled “Balancing Leasing and Insurance Costs to Achieve Total Risk Coverage in Cloud Storage Multi-Homing” focused on characterizing “insurance” costs due to unavailability of a Cloud provider [9]. The author developed a model describing how a penalty value could be associated with unavailability. The contribution by Galati et al. entitled “A WS-Agreement Based SLA Implementation for the CMAC Platform” described how SLAs could be negotiated and monitored based on the “Condition Monitoring on a Cloud” platform [10]. The authors demonstrated their approach using the Cybula engine for large-scale data analysis. The last paper in this session, “A Domain Specific Language and a Pertinent Business Vocabulary for Cloud Service Selection” by Slawik and Küpper [11], described a domain-specific language for describing services and a business vocabulary for supporting selection and brokering of these services. The authors suggest that as additional providers become available, such a vocabulary is necessary to enable comparison across providers.

The final session focused on economic aspects and modeling strategies for supporting quality of service. The first contribution by Tolosana-Calasanz, Bañares, and Colom entitled “Towards Petri Net-Based Economical Analysis for Streaming Applications Executed over Cloud Infrastructures” described how an application deployed over a distributed computing infrastructure (focusing, in this case, on a systolic array-based architecture) could be modeled from both a “functional” and “operational” perspective [12]. The authors used Reference Nets (a type of Petri net) for developing

the model and analyzing it with reference to streaming applications. The contribution by Butoi, Stan, and Silaghi entitled “Autonomous Management of Virtual Machine Failures in IaaS Using Fault Tree Analysis” described strategies for supporting VM migration based on predicted failures in the Cloud infrastructure (which can have varying availability and reliability characteristics) [13]. The authors made use of Xen-based fault traces to demonstrate how their approach, based on fault trees, could be used in practice. A key contribution was to determine how a node within such an infrastructure could accept future jobs or delegate jobs to other nodes based on likely availability profiles. The last paper in this session, entitled “How Do Content Delivery Networks Affect the Economy of the Internet and the Network Neutrality Debate?” by Maillé and Tuffin [14], described how economic models could be associated with content delivery networks involving a combination of content providers and Internet service providers (ISPs). The authors described how content provisioning could be shared across multiple ISPs and the associated potential revenue models based on payment by a sender or receiver of the content. The authors also investigated how content should be cached across different providers and how caching benefits ISPs.

In addition to these sessions, we included the paper by Massimo Felici. This paper, which is aligned with Felici’s keynote at GECON 2014 [15], discusses the economics of security and investigates how economics may drive operational security and the deployment of security technologies.

Finally, we would like to wholeheartedly thank the reviewers and Program Committee members for completing their reviews on time, and giving insightful and valuable feedback to the authors. Furthermore, we would like to thank Alfred Hofmann of Springer for his support in publishing the proceedings of GECON 2014. The collaboration with Alfred Hofmann and his team has been, as in the past years, efficient and effective.

September 2014

Kurt Vanmechelen
Jörn Altmann
Omer F. Rana

References

1. Kauffman, R.J., Ma, D., Yu, M.: A metrics suite for firm-level cloud computing adoption readiness. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 19–35. Springer, Heidelberg (2014)
2. Pallas, F.: An agency perspective to cloud computing. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 36–51. Springer, Heidelberg (2014)
3. Kim, K., Kang, S., Altmann, J.: Cloud goliath versus a federation of cloud dauids: Survey of economic theories on cloud federation. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 55–66. Springer, Heidelberg (2014)
4. Vega, D., Meseguer, R., Freitag, F.: Analysis of the social effort in multiplex participatory networks. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 67–79. Springer, Heidelberg (2014)
5. Lučanin, D., Jrad, F., Brandic, I., Streit, A.: Energy-aware cloud management through progressive sla specification. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 83–98. Springer, Heidelberg (2014)

6. Douglas, G., Drawert, B., Krantz, C., Wolski, R.: CloudTracker: using execution provenance to optimize the cost of cloud use. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 99–113. Springer, Heidelberg (2014)
7. Me, G.: Migration to governmental cloud digital forensics community: economics and methodology. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 114–129. Springer, Heidelberg (2014)
8. O’Loughlin, J., Gillam, L.: Performance evaluation for cost-efficient public infrastructure cloud use. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 133–145. Springer, Heidelberg (2014)
9. Naldi, M.: Balancing leasing and insurance costs to achieve total risk coverage in cloud storage multi-homing. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 146–158. Springer, Heidelberg (2014)
10. Galati, A., Djemame, K., Fletcher, M., Jessop, M., Weeks, M., McAvoy, J.: A WS-Agreement based SLA implementation for the CMAC platform. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 159–171. Springer, Heidelberg (2014)
11. Slawik, M., Küpper, A.: A domain specific language and a pertinent business vocabulary for cloud service selection. In: Altmann, J., Vanmechelen, K., Rana, O.F., (eds.) GECON 2014. LNCS, vol. 8914, pp. 172–185. Springer, Heidelberg (2014)
12. Tolosana-Calasanz, R., Bañares, J.A., Colom, J.M.: Towards petri net-based economical analysis for streaming applications executed over cloud infrastructures. In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 189–205. Springer, Heidelberg (2014)
13. Butoi, A., Stan, A., Silaghi, G.C.: Autonomous management of virtual machine failures in iaas using fault tree analysis. In: Altmann, J., Vanmechelen, K., Rana, O.F., (eds.) LNCS, vol. 8914, pp. 206–221. Springer, Heidelberg (2014)
14. Maillé, P., Tuffin, B.: How do content delivery networks affect the economy of the internet and the network neutrality debate? In: Altmann, J., Vanmechelen, K., Rana, O.F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 222–230. Springer, Heidelberg (2014)
15. Felici, M.: Economics, security and innovation. In: Altmann, J., Vanmechelen, K., Rana, O. F. (eds.) GECON 2014. LNCS, vol. 8914, pp. 3–15. Springer, Heidelberg (2014)



<http://www.springer.com/978-3-319-14608-9>

Economics of Grids, Clouds, Systems, and Services
11th International Conference, GECON 2014, Cardiff,
UK, September 16-18, 2014. Revised Selected Papers.
Altmann, J.; Vanmechelen, K.; Rana, O.F. (Eds.)
2014, XIV, 231 p. 63 illus., Softcover
ISBN: 978-3-319-14608-9