

## Theme section on performance modelling and engineering of software and systems

Catalina M. Lladó<sup>1</sup> · Kai Sachs<sup>1</sup>

Received: 7 August 2017 / Accepted: 4 September 2017 / Published online: 27 September 2017  
© Springer-Verlag GmbH Germany 2017

The modelling of the performance of today's computer systems is getting more and more challenging. While most models traditionally focused on utilization and response time of a single system, technologies such as cloud and containerization require new approaches to efficiently handle the complexity of current installations. The meaning of the word performance is constantly changing—currently covers additional non-functional aspects including security and reliability. In a world full of software systems composed of huge number of microservices running on cloud infrastructures, a single service is less important than in the traditional system architectures. In contrast to the classical systems, the entire network of containerized services needs to be efficiently operated, monitored and orchestrated. Containers are automatically deployed, scaled and operated by orchestration middleware, IaaS is re-arranging their resources dynamically, and several independent services might compete for the same resources. To optimize efficiency users, developers and operators need new approaches to model systems and their behaviour, allowing not only to evaluate the systems retrospectively, but also to take the right actions proactively in near real time.

At the International Conference on Performance Engineering (ICPE) 2015 in Austin, we discussed the meaning of performance modelling at both, the main conference and the various co-located workshops. In particular, the participants of the *First Workshop on Challenges in Performance Methods for Software Development* were very excited to discuss and work on this topic. Due to the importance and

high interest in this topic, we decided to focus this theme section on the performance modelling aspects and related challenges. Overall, we selected six high-quality submissions to reflect different aspects such as SLAs, scalability of cloud systems and workload specification. The first paper, *Accelerating task completion in mobile offloading systems through adaptive restart* by Wang and Wolter, presents an upgrade of the restart mechanism used to mitigate the effects of packet loss and delay in wireless networks on the completion time of an offloading task. Following, the paper *Model-driven performance prediction of systems of systems* by Falkner, Szabo, Chiprianov, Puddy, Rieckmann, Fraser and Aston proposes an integrated approach to performance prediction of model-driven real-time embedded defence systems and systems of systems. Next, Vögele, van Hoorn, Schulz, Hasselbring and Krcmar present an approach that aims at automating the extraction and transformation of workload specifications for load testing and model-based performance prediction. Their paper is entitled *WESSBAS: extraction of probabilistic workload specifications for load testing and performance prediction—a model-driven approach for session-based application systems*.

The paper *Using UML/MARTE to support performance tuning and stress testing in real-time systems* by Di Alessio and Sen provides an approach, to support the generation of system configurations predicted to achieve a satisfactory trade-off between response time and CPU usage. In the paper, the authors also provide stress test cases that push the system tasks to violate their deadlines. The next paper, *Model-driven optimal resource scaling in cloud* by Gandhi, Dube, Karve, Kochut and Zhang, presents a solution for optimizing the resource scaling of cloud deployments via implementation in OpenStack as the underlying cloud platform. Finally, Buchholz and Vastag—in their paper *Towards an Analytical Method for SLA Validation*—propose an analytical approach

---

✉ Kai Sachs  
kai.sachs@sap.com

Catalina M. Lladó  
cllado@uib.cat

<sup>1</sup> Gainesville, FL, USA

for the analysis of service-oriented architectures based on the information available in service-level agreements.

**Acknowledgements** We would like to thank Martin Schindler for his great support and the authors for making this theme section possible. We are also thankful to the reviewers for their detailed reviews which in many cases evolve to substantial improvements. Further, we would like to thank the entire ICPE community for their support and numerous fruitful discussions.



**Catalina M. Lladó** is an Associate Professor at the Maths and Computer Science Department of the Universitat de les Illes Balears. Her research interests include, but are not limited to, performance modelling of computing systems, software performance engineering and performance interchange formats. She received her Ph.D. in Computer Science from Imperial College, London, in 2002.



**Kai Sachs** is the Development Manager of the Healthcare Networks team at the SAP Health Innovation Hub in Potsdam. In his previous role as Principal Development Expert, he was responsible for strategic projects in the Healthcare Development Unit. Before moving to the healthcare domain, he was focusing on large-scale systems, including cloud computing and big data technologies. He received a Ph.D. degree from Technische Universität Darmstadt, which was awarded with the SPEC Distinguished Dissertation Award for outstanding contributions.