Special Issue: Selected Papers From the 8th ACM/SPEC International Conference on Performance Engineering (ICPE 2017)

ANNE KOZIOLEK, Karlsruhe Institute of Technology
 EVGENIA SMIRNI, College of William and Mary

We are delighted to present in this special issue of TOMPECS three of the highest-rated papers of the Research Track of the 8th ACM/SPEC International Conference on Performance Engineering (ICPE) that took place in L'Aquila, Italy, April 22–26, 2017. The goal of ICPE is to integrate theory and practice in the field of performance engineering by providing a forum for sharing ideas and experiences between industry and academia. The conference grew out of the ACM Workshop on Software Performance (WOSP since 1998) and the SPEC International Performance Engineering Workshop (SPEW since 2008).

In the ICPE 2017 edition, the research track attracted 65 submissions with 22 accepted for presentation as full papers. The program covered traditional ICPE topics such as design for performance and problem diagnosis, online performance management, analytic models, empirical studies, model building, and benchmarking, as well as application of performance engineering theory and techniques to several practical fields, including distributed systems, cloud computing, storage, energy, big data, virtualized systems, and hardware.

In this special issue, we present the extended versions of three of the highest-rated papers in the conference. Three members of the ICPE 2017 program committee served as reviewers for each extended version, and each paper went through the regular TOMPECS review/revision cycle.

The first article by Cheng Wang (VMware Inc.), Qianlin Liang (Ele.me Inc.), and Bhuvan Urgaonkar (Penn State University), titled “An Empirical Analysis of Amazon EC2 Spot Instance Features Affecting Cost-Effective Resource Procurement,” focuses on ways to optimize the usage of Amazon’s popular EC2 spot markets. The article defines four features that can result in a better tenant bid: instance lifetime, average spot price during lifetime, simultaneous revocations, and startup delay. For these features, they illustrate how current prediction practices are inadequate and propose new and more accurate prediction alternatives. The proposed predictors are extensively evaluated using real-world workloads and show that their approach results in better performance (measured in diminished bid failures and lower tail latency) with the same cost as prior work. Results are corroborated with two sets of trace data: a 3-month period in 2015 and another 3-month period in 2017. The work is a case study in traditional ICPE topics: problem diagnosis and performance prediction in cloud computing.

Authors’ addresses: A. Koziolek, Karlsruhe Institute of Technology, Am Fasanengarten 5, Karlsruhe, Germany; email: koziolek@kit.edu; E. Smirni, College of William and Mary, 104 Jamestown Rd., Williamsburg, VA, 23185; email: esmirni@cs.wm.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

2018 Copyright is held by the owner/author(s).

2376-3639/2018/03-ART6e

<https://doi.org/10.1145/3186329>

The second article by Benjamin Cassell, Tyler Szepesi, Jim Summers, Tim Brecht (all from the University of Waterloo), Derek Eager (University of Saskatchewan), and Bernard Wong (University of Waterloo), titled “Disk Prefetching Mechanisms for Increasing HTTP Streaming Video Server Throughput,” describes Libception, a portable system in which techniques for optimizing IO in HTTP video streaming can be implemented and evaluated. The authors show that combining disk access serialization and aggressive prefetching of requests is fundamental to improve user-perceived performance. Libception implements these techniques seamlessly. A Web server that is linked with Libception can achieve the benefits of serialization and prefetching, regardless of which disk scheduler is used, essentially achieving nearly as good performance as kernel parameter tuning. The authors present extensive evidence of Libception’s ability to improve video streaming performance using a variety of Web servers on multiple disks and additionally show the potential of Libception for improving other workloads beyond streaming. This article again focuses on problem diagnosis and the development of a systems solution to improve performance.

The third article by Alexey Ilyushkin (Delft University of Technology), Ahmed Ali-Eldin (Umeå University and University of Massachusetts at Amherst), Nikolas Herbst and André Bauert (University of Würzburg), Alessandro V. Papadopoulos (Mälardalen University), Dick Epema (Delft University of Technology), and Alexandru Iosup (Vrije Universiteit Amsterdam and Delft University of Technology), titled “An Experimental Performance Evaluation of Autoscalers for Complex Workflows,” focuses on a classic cloud performance issue: the effect of autoscalers on workload performance. The authors focus on workflow workloads that can be best characterized as a set of tasks with precedence constraints among themselves. They perform a comprehensive experimental evaluation (across 15 performance metrics) of several autoscalers that have appeared in the literature, along with two autoscalers that are proposed here and are tailored specifically for workflow workloads. This work addresses questions such as how different autoscalers affect deadline violations, how they affect costs, and how they find a balance between performance and cost.

A common characteristic of the three selected papers is that they are all heavily experimental and represent excellent examples of practical performance evaluation in the real world.

With this special issue, it is our hope that the readers of TOMPECS will get a sample of the research work that ICPE publishes.