Introduction to special issue: ACM SIGMETRICS 2016

Thomas Bonald[1] · **Ayalvadi Ganesh**[2]

The four papers in this special issue are based on a selection of top theoretical papers from the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, which took place in Antibes Juan-les-Pins, France, in June 2016. As the premier international conference on the measurement and modeling of computer systems, the technical program featured papers that covered both theory and applications from a wide variety of areas. We seek with this special issue to highlight some of the recent theoretical work on the mathematical analysis and modeling of computer systems.

The paper of Ferragut, Rodriguez and Paganini studies timer-based caching policies. It is known for Poisson inter-request times (with rates dependent on item popularity) that the optimal policy is to cache the most popular items. The paper shows that this result extends to inter-request times with increasing hazard rates, but not to decreasing hazard rates. In the latter case, the problem is shown to reduce to a convex program, and a solution presented.

Jonckheere and Prabhu study a load balancing problem with Poisson arrivals, general service times, and identical servers employing an insensitive service discipline such as processor sharing, which leads to queue length distributions that depend only on the mean service time. They study asymptotics in the Halfin-Whitt scaling regime, where both the arrival rate and the number of servers tend to infinity, in such a way that the system is in heavy traffic. They obtain different scalings for the queue length in subcritical, critical and supercritical loading regimes.

✉ Ayalvadi Ganesh
   A.Ganesh@bristol.ac.uk

[1]   Telecom Paris-Tech, Paris, France

[2]   University of Bristol, Bristol, UK

The paper of Maguluri, Burle and Srikant considers the heavy-traffic regime of an input switch where only a subset of ports are saturated. It shows that max-weight scheduling is optimal in terms of total queue length in the limit of a large number of ports. Interestingly, the ports are allowed to saturate at different rates. The drift technique used in the analysis is generic and applicable to a broad class of systems.

Hybrid switching combining circuit switching and packet switching is the topic of the paper of Venkatakrishnan, Alizadeh and Viswanath. The salient feature of hybrid switching is the cost of reconfiguration of the high-bandwidth circuit switch. The authors propose a nearly optimal scheduling algorithm that triggers reconfigurations when needed. The analysis relies on the submodular structure of the underlying optimization problem.

We are grateful to the authors for amending and expanding their conference papers into journal versions tailored to the Queueing Systems community. The papers have gone through a regular Queueing Systems review cycle, and we are indebted to the anonymous referees not only for the quality and thoroughness of their work.