Editorial introduction

A. B. Dieker · M. Lelarge

Published online: 9 May 2014
© Springer Science+Business Media New York 2014

The five 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 at Carnegie Mellon University, Pittsburgh, PA, in June 2013. As the premier 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, including algorithms, communication networks, dynamics and control, optimization, performance analysis, resource allocation and scheduling, and stochastic modeling, among others. The program consisted of 26 papers selected from nearly 200 international submissions by a technical program committee of well-established researchers from all over the world.

With this special issue, we seek to highlight some of the recent theoretical work on the mathematical analysis and modeling of computer systems. These papers are motivated by applications related to computers and use rigorous mathematical methods to study these problems. The five papers presented herein represent a small fraction of the theoretical work presented at the 2013 ACM SIGMETRICS Conference; refer to <http://www.sigmetrics.org/sigmetrics2013> for additional details.

The papers of Li, Lee, and Lui and of Van Houdt use stochastic models to study input/output dynamics for solid-state drives (SSDs), aiming at the efficient design of garbage collection algorithms. In SSDs, data are stored in pages, which in turn are organized in blocks. Since only blocks can be erased as opposed to individual pages,

A. B. Dieker
H. Milton Stewart School of Industrial and Systems Engineering,
Georgia Institute of Technology, Atlanta, GA30332, USA
e-mail: ton.dieker@isye.gatech.edu

M. Lelarge (✉)
INRIA-Ecole Normale Supérieure, Paris, France
e-mail: marc.lelarge@ens.fr

blocks must periodically be cleaned, and this task is fulfilled by garbage collection algorithms. Different algorithms have different effects on speed and wear. These papers use mean-field techniques to study the performances of several garbage collection algorithms.

The paper of Gandhi, Doroudi, Harchol-Balter, and Scheller-Wolf is motivated by the desire to reduce power consumption of data centers while at the same time maintaining service and availability levels. It focuses on multiserver queueing systems with setup times, meaning that an idling server can only be made available for service after some (random) setup time. The paper also studies variations where idle servers delay for a period of time before turning off or can be put to sleep.

The papers of Bouman, Borst, and Van Leeuwen and of Simatos, Bouman, and Borst are motivated by questions on queueing delay for wireless networks operating under distributed queue-based random-access schemes. Much existing work predominantly focuses on throughput, but the study of delays calls for other, finer techniques. Bouman et al. establishes lower bounds for the average steady-state delay, which offer insight in the cause of the excessive delays. Simatos et al. focuses on the effects of activation rules on delays and establishes a peculiar ‘lingering’ phenomenon for a class of networks.

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 of their work, but also for their assistance in working with our tight timeline. We thank Sergey Foss for inviting us to edit this special issue and for his support throughout the process. We also thank the editorial staff of Queueing Systems for their help in preparing this special issue.