Queueing Systems (2019) 91:205–206
https://doi.org/10.1007/s11134-019-09607-0

Introduction to special issue: The IFIP Performance 2018 conference

Michel Mandjes[1] · Philippe Robert[2]

Published online: 13 March 2019

This volume of "Queueing Systems: Theory and Applications" contains two papers that have been presented at the Performance 2018 Conference, or officially "The 36th International Symposium on Computer Performance, Modeling, Measurements and Evaluation 2018," held in Toulouse.

The conference is one of the top venues for dissemination of high-quality research in the area of performance analysis and quantitative assessment of communication networks. It has followed an 18-month cycle for many years. Recently, the frequency has been adapted to one conference per year.

The organization of the event has been done in close interaction with working group 7.3 of the International Federation for Information Processing (IFIP), focusing on Computer System Modelling. WG 7.3 is chaired by Mark Squillante of IBM/Yorktown Heights, as such also member of Perfomance's Steering Group, together with Sem Borst and Benny Van Houdt. The general chair of the conference was Urtzi Ayesta, while the chairs of the Technical Programme Committee were John Liu (Chinese University of Hong Kong), Michel Mandjes (University of Amsterdam), and Debasis Mitra (Columbia University). When setting up the call, special attention was paid to covering a timely set of topics. Objects of study include blockchain, smart grids, and net neutrality, next to the more classical topics. Also at the methodological level, emphasis was put on new developments, such as those in artificial intelligence and machine learning.

All submissions to the Performance conference have undergone a careful review procedure. They were evaluated through a double-blind process that ensures anonymity of both authors and reviewers. Each paper received at least three reviews from program committee members who are not in conflict of interest with the authors. The accepted papers had the option to be submitted to special issues of a number of journals, including Queueing Systems.

✉ Michel Mandjes
m.r.h.mandjes@uva.nl

1 University of Amsterdam, Amsterdam, The Netherlands

2 INRIA, Paris, France

The first contribution is "Delay Asymptotics and Bounds for Multi-Task Parallel Jobs" by Weina Wang, Mor Harchol-Balter, Haotian Jiang, Alan Scheller-Wolf, and R. Srikant. This paper studies an M/G/n queueing system, where each arriving task triggers sub-tasks in k^n queues (uniformly at random); the service times of sub-tasks are i.i.d. In this model, each task is completed only when all of the sub-tasks are completed, entailing that the delay of a task is the maximum of the delays of its sub-tasks. A natural approximation consists of assuming that each of the queues behave independently. When k^n is fixed this is a classical mean-field convergence result (when n goes to infinity). Little is known when k^n grows with respect to n. The two main results of the paper are:

1. When $k^n = o(n^{1/4})$ there is asymptotic independence (in a certain sense, made precise using the concept of total variation distance).
2. In the non-asymptotic regime, it is shown that the task delay is stochastically bounded by that in the corresponding system with independent queues.

The second contribution is "Temporal starvation in multi-channel CSMA networks: an analytical framework" by Alessandro Zocca. This paper studies a stochastic model for a frequency-agile CSMA protocol for wireless networks where multiple orthogonal frequency channels are available. The focus is on the asymptotic regime in which the network nodes try to activate aggressively so as to achieve maximum throughput. Special attention is paid to the scenario where the number of available channels is not sufficient for all nodes of the network to be simultaneously active, so that the well-studied temporal starvation issues of the single-channel CSMA dynamics persist. Perhaps surprisingly, the author proves that the aggregate throughput is a non-increasing function of the number of available channels. To further study the trade-off between aggregate throughput and temporal starvation phenomena, an analytical framework is used that captures the transient dynamics of multi-channel CSMA networks by means of first hitting times.