



Stochastic Systems

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Applied Probability Society Student Paper Competition: Abstracts of 2019 Finalists

To cite this article:

(2021) Applied Probability Society Student Paper Competition: Abstracts of 2019 Finalists. *Stochastic Systems* 11(4):391-393. <https://doi.org/10.1287/stsy.2021.0073>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Stochastic Systems*. Copyright © 2021 The Author(s). <https://doi.org/10.1287/stsy.2021.0073>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2021 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>


Applied Probability Society Student Paper Competition: Abstracts of 2019 Finalists

<https://doi.org/10.1287/stsy.2021.0073>

Copyright: © 2021 The Author(s)

The journal is pleased to publish the abstracts of the winner and finalists of the 2019 Applied Probability Society's student paper competition.

The 2019 student paper prize committee was chaired by Amy Ward. The 2019 committee members are (in alphabetical order by last name): Reza Aghajani, Pelin Canbolat, Jing Dong, Johan van Leeuwen, Ilya Ryzhov, Assaf Zeevi, Jiheng Zhang, and Serhan Ziya.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "Stochastic Systems. Copyright © 2021 The Author(s). <https://doi.org/10.1287/stsy.2021.0073>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>."

The 2019 prize winners are as follows:

First Prize

Sub-Diffusive Load-Balancing in Time-Varying Queueing Systems
Gal Mendelson, Technion–Israel Institute of Technology

Finalists (in alphabetical order according to student's last name):

Phase Transitions and Cyclic Phenomena in Bandits with Switching Constraints
Yunzong Xu, Massachusetts Institute of Technology

Marrying Stochastic Gradient Descent With Bandits: Learning Algorithms for Inventory Systems with Fixed Costs
Hao Yuan, University of Michigan

A Lower Bound on the Queueing Delay in Resource Constrained Load Balancing
Martin Zubeldia, Eindhoven University of Technology and the University of Amsterdam

Sub-Diffusive Load-Balancing in Time-Varying Queueing Systems

Gal Mendelson

Technion–Israel Institute of Technology, galmen@campus.technion.ac.il

Coauthors: Rami Atar, Technion–Israel Institute of Technology; Isaac Keslassy, Technion–Israel Institute of Technology

Load-balancing algorithms for systems that operate in heavy traffic are known to lead, under suitable conditions, to state space collapse (SSC). This term refers to the phenomenon where imbalance is negligible compared with queue lengths. Specifically, whereas queue lengths behave diffusively, the size of imbalance is at a subdiffusive scale: denoting by n the usual scaling parameter, the former and the latter are of order $O(n^{\frac{1}{2}})$ and $o(n^{\frac{1}{2}})$, respectively. In this paper, we consider load balancing for time-varying systems. SSC results and the standard techniques on which they are based do not apply to these systems, which (a) are not in heavy traffic, thus queue lengths may reach levels as high as $O(n)$, and (b) have time-varying traffic intensities that cause transitions between underloaded, critically loaded, and overloaded regimes. Our results extend SSC far beyond the heavy traffic setting, by establishing subdiffusive (i.e., $o(n^{\frac{1}{2}})$) balance for time-varying systems.

To exhibit the breadth of the described phenomenon, the results address three load-balancing models. The first is *the-power-of- d -choices* ($SQ(d)$), where arrivals from a single stream are routed to the shortest among d

randomly chosen queues, where $1 < d \leq N$, and N denotes the fixed number of queues in the system. The second is *redundancy- d* ($R(d)$), where jobs are replicated d times, routed simultaneously to d randomly chosen queues, and all but the first replica to be admitted into service are canceled. The third model is *longest queue first* (LQF), where a single resource is shared by N job classes, and the job that receives service is always selected from the queue that is longest.

As an application of these results, asymptotic optimality of $SQ(d)$ and $R(d)$ is shown, with an optimality guarantee of order $o(n^{\frac{1}{2}})$ in the aforementioned framework, where in particular queue sizes may reach $O(n)$. Moreover, in the special case of the standard heavy traffic setting, the results are shown to yield new, explicit sufficient conditions for SSC.

Phase Transitions and Cyclic Phenomena in Bandits with Switching Constraints

Yunzong Xu

Massachusetts Institute of Technology, yxu@mit.edu

Coauthor: David Simchi-Levi, Massachusetts Institute of Technology

We consider the classical stochastic multiarmed bandit problem with a constraint on the total cost incurred by switching between actions. We prove matching upper and lower bounds on regret and provide near-optimal algorithms for this problem. Surprisingly, we discover phase transitions and cyclic phenomena of the optimal regret. That is, we show that associated with the multiarmed bandit problem, there are phases defined by the number of arms and switching costs, where the regret upper and lower bounds in each phase remain the same and drop significantly between phases. The results enable us to fully characterize the trade-off between regret and incurred switching cost in the stochastic multiarmed bandit problem, contributing new insights to this fundamental problem. Under the general switching cost structure, the results reveal a deep connection between bandit problems and graph traversal problems, such as the shortest Hamiltonian path problem.

Marrying Stochastic Gradient Descent With Bandits: Learning Algorithms for Inventory Systems with Fixed Costs

Hao Yuan

University of Michigan, haoyuan@umich.edu

Coauthors: Qi Luo, Clemson University; Cong Shi, University of Michigan

We consider a periodic-review single-product inventory system with fixed cost under censored demand. Under full demand distributional information, it is well-known that the celebrated (s, S) policy is optimal. In this paper, we assume the firm does not know the demand distribution a priori, and makes adaptive inventory ordering decision in each period based only on the past sales (that is, censored demand) data. The standard performance measure is regret, which is the cost difference between a feasible learning algorithm and the clairvoyant (full-information) benchmark. Compared with prior literature, the key difficulty of this problem lies in the loss of joint convexity of the objective function, due to the presence of fixed cost. We develop a nonparametric learning algorithm termed the (δ, S) policy that combines the powers of stochastic gradient descent, bandit controls, and simulation-based methods in a seamless and nontrivial fashion. We prove that the cumulative regret is $O(\log T \sqrt{T})$, which is provably tight up to a logarithmic factor. We also develop several technical results that are of independent interest. We believe that the framework developed could be widely applied to learning other important stochastic systems with partial convexity in the objectives.

A Lower Bound on the Queuing Delay in Resource Constrained Load Balancing

Martin Zubeldia

Eindhoven University of Technology and the University of Amsterdam, m.zubeldia.suarez@tue.nl

Coauthors: David Gamarnik, Massachusetts Institute of Technology; John N. Tsitsiklis, Massachusetts Institute of Technology

We consider the following distributed service model: jobs with unit mean, general distribution, and independent processing times arrive as a renewal process of rate λn , with $0 < \lambda < 1$, and are immediately dispatched to one of several queues associated with n identical servers with unit processing rate. We assume that

the dispatching decisions are made by a central dispatcher endowed with a finite memory, and with the ability to exchange messages with the servers. We study the fundamental resource requirements (memory bits and message exchange rate), in order to drive the expected queueing delay in steady-state of a typical job to zero, as n increases. We develop a novel approach to show that, within a certain broad class of “symmetric” policies, every dispatching policy with a message rate of the order of n , and with a memory of the order of $\log n$ bits, results in an expected queueing delay that is bounded away from zero, uniformly as $n \rightarrow \infty$. This complements existing results that show that in the absence of such limitations on the memory or the message rate, there exist policies with vanishing queueing delay (at least with Poisson arrivals and exponential service times).