



Stochastic Systems

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

Applied Probability Society Student Paper Competition: Abstracts of 2018 Finalists

To cite this article:

(2019) Applied Probability Society Student Paper Competition: Abstracts of 2018 Finalists. *Stochastic Systems* 9(1):81–82. <https://doi.org/10.1287/stsy.2018.0025>

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 © 2019 The Author(s). <https://doi.org/10.1287/stsy.2018.0025>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2019, 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 2018 Finalists

<https://doi.org/10.1287/stsy.2018.0025>

Copyright: © 2019 The Author(s)

The journal is pleased to publish the abstracts of the winner and finalists of the 2018 Applied Probability Society's student paper competition.

The 2018 student paper prize committee was chaired by John Hasenbein. The 2018 committee members are (in alphabetical order by last name) Reza Aghajani, Rami Atar, Jose Blanchet, Pelin Canbolat, Jing Dong, Xin Guo, Andreea Minca, Petar Momcilovic, Amber Puha, Ilya Ryzhov, and Galit Yom-Tov.



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 © 2019 The Author(s). <https://doi.org/10.1287/stsy.2018.0025>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>."

The 2018 prize winners are as follows:

First Prize

Statistics of Robust Optimization: A Generalized Empirical Likelihood Approach

Hongseok Namkoong, Stanford University

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

Adaptive Learning with Unknown Information Flows

Ahmadreza Momeni, Stanford University

SOAP: One Clean Analysis of All Age-Based Scheduling Policies

Ziv Scully, Carnegie Mellon University

Stochastic Games for Fuel Followers Problem: N vs. MFG

Renyuan Xu, University of California, Berkeley

Statistics of Robust Optimization: A Generalized Empirical Likelihood Approach

Hongseok Namkoong

Stanford University, hnamk@stanford.edu

Advisors: John C. Duchi, Stanford University;

Peter W. Glynn, Stanford University

We study statistical inference and distributionally robust solution methods for stochastic optimization problems, focusing on confidence intervals for optimal values and solutions that achieve exact coverage asymptotically. We develop a generalized empirical likelihood framework—based on distributional uncertainty sets constructed from nonparametric f -divergence balls—for Hadamard differentiable functionals and, in particular, stochastic optimization problems. As consequences of this theory, we provide a principled method for choosing the size of distributional uncertainty regions to provide one- and two-sided confidence intervals that achieve exact coverage. We also give an asymptotic expansion for our distributionally robust formulation, showing how robustification regularizes problems by their variance. Finally, we show that optimizers of the distributionally robust formulations we study enjoy (essentially) the same consistency properties as those in classical sample average approximations. Our general approach applies to quickly mixing stationary sequences, including geometrically ergodic Harris recurrent Markov chains.

Adaptive Learning with Unknown Information Flows

Ahmadreza Momeni

Stanford University, amomenis@stanford.edu

Advisor: Yonatan Gur, Stanford University

An agent facing sequential decisions that are characterized by partial feedback needs to strike a balance between maximizing immediate payoffs based on available information and acquiring new information that may be

essential for maximizing future payoffs. This trade-off is captured by the multiarmed bandit (MAB) framework, in which at each time epoch a single observation is collected on the action that was selected at that epoch. However, in many practical settings, additional information may become available between pulls, which might be essential for achieving good performance. We introduce a generalized MAB formulation that relaxes the strong assumptions on the information-collection process and in which auxiliary information on each arm may appear arbitrarily over time. By obtaining matching lower and upper bounds, we characterize the (regret) complexity of this family of MAB problems as a function of the information flows. We introduce a broad adaptive exploration approach for designing policies that, without any prior knowledge on the information-arrival process, attain the best performance (regret rate) that is achievable when the information-arrival process is a priori known. Our approach is based on adjusting MAB policies designed to perform well in the absence of auxiliary information by using dynamically customized virtual time indexes to endogenously control the exploration rate of the policy. We demonstrate the effectiveness of our approach through establishing performance bounds and evaluating numerically the performance of adjusted well-known MAB policies. Our study demonstrates how decision-making policies designed to perform well with very little information can be adjusted to also guarantee optimality in more information-abundant settings.

SOAP: One Clean Analysis of All Age-Based Scheduling Policies

Ziv Scully

Carnegie Mellon University, zscully@cs.cmu.edu

Advisors: Mor Harchol-Balter, Carnegie Mellon University;

Alan Scheller-Wolf, Carnegie Mellon University

We consider an extremely broad class of $M/G/1$ scheduling policies called SOAP: schedule ordered by age-based priority. The SOAP policies include almost all scheduling policies in the literature as well as an infinite number of variants that have never been analyzed or maybe not even conceived. SOAP policies range from classic policies, such as first-come first-served, foreground-background, class-based priority, and shortest remaining processing time; to much more complicated scheduling rules, such as the famously complex Gittins index policy and other policies in which a job's priority changes arbitrarily with its age. Although the response time of policies in the former category is well understood, policies in the latter category have resisted response-time analysis. We present a universal analysis of all SOAP policies, deriving the mean and Laplace–Stieltjes transform of response time.

Stochastic Games for Fuel Followers Problem: N vs. MFG

Renyuan Xu

UC Berkeley, renyuanxu@berkeley.edu

Advisor: Xin Guo, University of California, Berkeley

In this paper, we formulate and analyze an N -player stochastic game of the classical fuel follower problem and its mean field game (MFG) counterpart. For the N -player game, we obtain the Nash equilibrium (NE) explicitly in two steps: first is to derive and analyze a system of Hamilton–Jacobi–Bellman equations, and second is to establish the existence of a unique strong solution to the associated Skorokhod problem on an unbounded polyhedron with oblique reflections. For the MFG, we derive a bang bang-type NE by a viscosity solution approach under some mild technical conditions. We also show that this solution is an ϵ -NE to the N -player game with $\epsilon = O(\frac{1}{\sqrt{N}})$. The N -player game and the MFG differ in that the NE for the former is state dependent and the NE for the latter is state independent. Our analysis shows that the NE for a stationary MFG may not be an NE for the corresponding nonstationary MFG.