INFORMS APS 2019

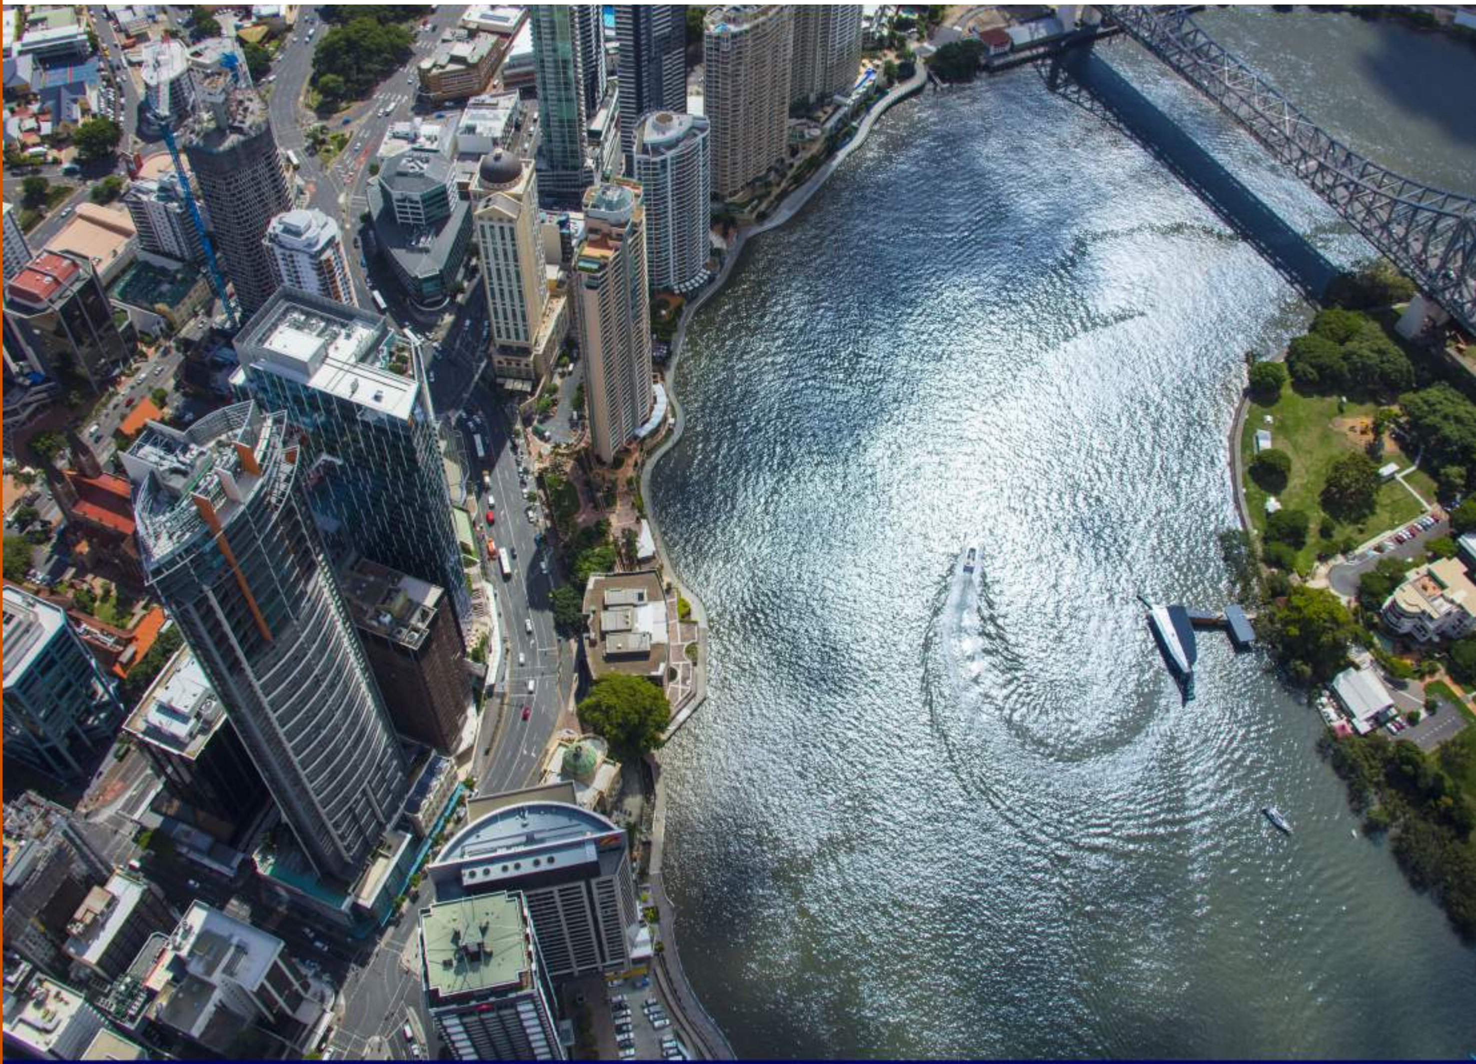

Brisbane
July 3 – 5, 2019

Contents

1

Welcome Address

Dear Participant,

It is with great pleasure that we welcome you to the 20th INFORMS Applied Probability Society conference in Brisbane, Australia. For many of you, it is not a trivial matter to cross the Pacific Ocean, all of Asia, or the Indian Ocean to arrive to Down Under. Still you are here, ready to exchange ideas and engage with colleagues from across the globe. Thank you for making the journey!

When it comes to Applied Probability, Australia and New Zealand have always been significant players. In the past, classic giants such as Daryl Daley, Joe Gani, Christopher Heyde, Eugene Seneta, Richard Tweedie, and David Vere-Jones have made major contributions to all areas of Applied Probability. Today, active researchers are present in almost every major city including Adelaide, Auckland (NZ), Brisbane, Canberra, Hobart, Melbourne, Perth, and Sydney. In addition, several research leaders in North America, Europe and Asia have come from this part of the world. Hence, we the Australian and New Zealand Applied Probability researchers feel very connected to the international Applied Probability community, and are delighted to see so many researchers at this conference.

This three-day conference hosts almost 280 talks, of which four are plenary lectures and two are tutorials. Plenary lectures are 50 minutes in duration, and tutorials are 90 minutes each. For the invited and contributed sessions, there are generally eight parallel sessions, each consisting of up to four 20-minute talks. The rich program features talks that deal with analysis of congestion systems, stochastic optimisation methods, learning theory, analysis of stochastic processes, simulation, optimal control, healthcare management, financial mathematics and many other aspects of Applied Probability. On the first day, the conference also has an evening event celebrating women in Applied Probability. On the calendar, the conference neighbours four additional events: the workshop in Honour of Peter Taylor dealing with Queues, Modelling and Markov Chains, the Applied^2 Probability Day, the Workshop on Algorithms and Models for the Web Graph (WAW 2019), and the 12th International Conference on Monte Carlo Methods and Applications (MCM2019).

This conference would not have been possible without the help of over 50 session organizers and of the conference organization team, the scientific advisory committee, the support of the APS Council, the guidance of previous APS Chair, Amy Ward, and current APS Chair, Itai Gurvich, and others. With the financial support of APS, nearly 30 students received funding. We also thank the University of Queensland's School of Mathematics and Physics, the ARC Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS), and the Applied Probability Special Interest Group of the Australian Mathematical Society (AustMS). Finally, we also thank PhD students for volunteering to help with registration and other aspects of the program.

With over 250 delegates, we are confident that these three days will yield fruitful scientific interactions. We hope that you also enjoy Brisbane and wish you a wonderful conference.

On behalf of the Organization Committee,

Yoni Nazarathy, the University of Queensland,
Giang Nguyen, the University of Adelaide.

2

Organization

Organizers

Yoni Nazarathy	The University of Queensland, Australia
Giang Nguyen	The University of Adelaide, Australia
Nigel Bean	The University of Adelaide, Australia
Marijn Jansen	The University of Queensland, Australia
Roxanne Jemison	The University of Queensland, Australia
Dirk Kroese	The University of Queensland, Australia
Ross McVinish	The University of Queensland, Australia
Sarat Babu Moka	The University of Queensland, Australia
Oscar Peralta	The University of Adelaide, Australia
Phil Pollett	The University of Queensland, Australia
Fred Roosta-Khorasani	The University of Queensland, Australia
Thomas Taimre	The University of Queensland, Australia
Peter Taylor	The University of Melbourne, Australia
Rixon Crane	The University of Queensland, Australia
Vektor Dewanto	The University of Queensland, Australia
Liam Hodgkinson	The University of Queensland, Australia
Yang Liu	The University of Queensland, Australia
Russell Tsuchida	The University of Queensland, Australia

Session Organizers

Ivo Adan	Eindhoven University of Technology, the Netherlands
Zeynep Akşin Karaesmen	Koç University, Turkey
Nilay Tanik Argon	University of North Carolina, USA
Alessandro Arlotto	Duke University, USA
Rami Atar	Technion – Israel Institute of Technology, Israel
Hayriye Ayhan	Georgia Institute of Technology, USA
François Baccelli	The University of Texas at Austin, USA
Chaithanya Bandi	Northwestern University, USA
Yonit Baron	Ariel University, Israel
Mohsen Bayati	Stanford University, USA
Sandjai Bhulai	VU University Amsterdam, the Netherlands
John Birge	University of Chicago, USA
Jose Blanchet	Columbia University, USA
Onno Boxma	Eindhoven University of Technology, the Netherlands
Peter Braunsteins	The University of Queensland, Australia
Amarjit Budhiraja	University of North Carolina, USA
Agostino Capponi	Columbia University, USA
Nan Chen	Chinese University of Hong Kong, Hong Kong
Asaf Cohen	University of Haifa, Israel

Jim Dai	Cornell University, USA
Duy-Minh Dang	The University of Queensland, Australia
Andrew Daw	Cornell University, USA
Jing Dong	Northwestern University, USA
Sherwin Doroudi	University of Minnesota, USA
Jan-Pieter Dorsman	University of Amsterdam, the Netherlands
Antonis Economou	University of Athens, Greece
Eugene Feinberg	Stony Brook University, USA
Sergey Foss	Heriot-Watt University, Scotland
David Goldberg	Cornell University, USA
Alexander Goldenshluger	University of Haifa, Israel
John Hasenbein	University of Texas at Austin, USA
Sophie Hautphenne	The University of Melbourne, Australia
Moshe Haviv	Hebrew University of Jerusalem, Israel
Shane Henderson	Cornell University, USA
Harsha Honnappa	Purdue University, USA
Susan Hunter	Purdue University, USA
Sandeep Juneja	Tata Institute of Fundamental Research, India
Jussi Keppo	National University of Singapore, Singapore
Yoav Kerner	Ben-Gurion University, Israel
Henry Lam	Columbia University, USA
Lingfei Li	Chinese University of Hong Kong, Hong Kong
Yunan Liu	North Carolina State University, USA
Siva Theja Maguluri	Georgia Institute of Technology, USA
Michel Mandjes	University of Amsterdam, the Netherlands
Pascal Moyal	Université de Lorraine, France
Sayan Mukherjee	Duke University, USA
Barry Nelson	Northwestern University, USA
Giang Nguyen	The University of Adelaide, Australia
Mariana Olvera-Cravioto	University of North Carolina at Chapel Hill, USA
Lerzan Örmeci	Koç University, Turkey
Guodong Pang	Pennsylvania State University, USA
Raghu Pasupathy	Purdue University, USA
Oscar Peralta	The University of Adelaide, Australia
Liron Ravner	University of Amsterdam, the Netherlands
Marty Reiman	Columbia University, USA
Chang-Han Rhee	Northwestern University, USA
Ad Ridder	VU University Amsterdam, the Netherlands
Fred Roosta-Khorasani	The University of Queensland, Australia
Dan Russo	Columbia University, USA
Pengyi Shi	Purdue University, USA
Flora Spieksma	Leiden University, the Netherlands
Mark Squillante	IBM Research – Thomas J. Watson Research Center, USA
Alexander Stolyar	University of Illinois at Urbana-Champaign, USA
Wen Sun	University of Southern California, USA
Baris Tan	Koç University, Turkey
Neil Walton	University of Manchester, England
Qiong Wang	University of Illinois at Urbana-Champaign, USA
Amy Ward	University of Chicago, USA
Yehua Wei	Boston College, USA
Adam Wierman	California Institute of Technology, USA
David Yao	Columbia University, USA
Nan Ye	The University of Queensland, Australia
Jiheng Zhang	Hong Kong University of Science and Technology, Hong Kong
Yuan Zhong	University of Chicago, USA
Enlu Zhou	Georgia Institute of Technology, USA
Alessandro Zocca	California Institute of Technology, USA

3

Sponsors

The organizers gratefully acknowledge the support of the following organizations:

https://connect.informs.org/aps/home

The University of Queensland
School of Mathematics and Physics
https://smp.uq.edu.au/

Centre for Applications in
Natural Resource Mathematics
https://smp.uq.edu.au/research/centres/carm/

http://www.acems.org.au/

http://www.austms.org.au/

4

Program at a Glance

Wednesday July 3, 2019				
	Track 1	Track 2	Track 3	Track 4
	Rooms P1 and P2	Room P3	Room P4	Room P5
8:00 –	**Registration**			
9:00 – 10:15	**Welcome & Keynote #1 (Sidney Resnick)** (Rooms P1 and P2)			
10:15 – 10:45	**Morning Tea**			
Session 1 *10:45 – 12:15*	*Stochastic Simulation Optimization* **Eunhye Song** Song Pasupathy Feng Nelson	*Stochastic Networks and Scaling Limits* **Amarjit Budhiraja** Atar Banerjee Puha Williams	*Job Scheduling and Load Balancing in Queueing Systems* **Onno Boxma & Jan-Pieter Dorsman** Cardinaels Mendelson Raaijmakers Dorsman	*Stochastic Systems* **Mark Squillante & David Yao** Taylor Agarwal Yao Feinberg
12:15 – 13:15	**Lunch**			
Session 2 *13:15 – 14:45*	*Model Uncertainty and Simulation Optimization* **Enlu Zhou** Song Qian Gao Zhou	*Queueing Networks and Markov Chains* **Jim Dai** Cao Choi Miyazawa	*Strategic Behavior in Queues* **Moshe Haviv** Kang Oz Wang Haviv	*Asymptotic Analysis in Stochastic Systems* **Guodong Pang** Borovkov Liu You Castiel
14:45 – 15:15	**Afternoon Tea**			
15:15 – 16:15	**Keynote #2 (Nelly Litvak)** (Rooms P1 and P2)			
16:15 – 16:30	**Mini Break**			
Session 3 *16:30 – 18:00*	*Efficiency in Stochastic Optimization and Simulation Experiments* **Henry Lam** Murthy Rhee Honnappa Qian	*Dynamic Investment and Stochastic Networks* **Agostino Capponi** Ji Sun Chen Lagos	*Level-Dependent QBD Processes* **Yonit Barron** Cordeiro Margolius Chen Barron	*Markov Chains* **Ivo Adan** Kapodistria Wickeham Dendievel Adan
18:30 – 20:00	**Celebrating Women in APS** — At the Ship Inn (15-Minute Walk)			

Wednesday July 3, 2019

	Track 5	Track 6	Track 7	Track 8
	Room M1	Room M2	Rooms M5 and M6	Rooms M7 and M8
8:00 –	**Registration**			
9:00 – 10:15	**Welcome & Keynote #1 (Sidney Resnick)** (Rooms P1 and P2)			
10:15 – 10:45	**Morning Tea**			
Session 1 *10:45 – 12:15*	*Statistical Methods for Learning and Decision-Making* **Nan Ye** Hu Liu Wu Snoswell	*Rare Events and Distributionally Robust Optimization* **Chang-Han Rhee** Blanchet Huang Su Rhee	*Control and Simulation of Stochastic Systems* **Jiheng Zhang** Chen Huang Zhang Frazier	*Stochastic Operations Research for Power Grids Applications* **Alessandro Zocca** Li Martyr Ferragut Zocca
12:15 – 13:15	**Lunch**			
Session 2 *13:15 – 14:45*	*Learning and Bandits* **Sandeep Juneja** Russo Yu Sundaresan Juneja	*Approximating Rare Event Probabilities* **Offer Kella** Hamilton Botev Bisewski Kella	*Asymptotics and Control* **Asaf Cohen** Ganesh Lam Lipshutz Cohen	*Revenue, Service, and Inventory Management* **Qiong Wang & Marty Reiman** Hu Huang Wang
14:45 – 15:15	**Afternoon Tea**			
15:15 – 16:15	**Keynote #2 (Nelly Litvak)** (Rooms P1 and P2)			
16:15 – 16:30	**Mini Break**			
Session 3 *16:30 – 18:00*	*Statistical and Computational Aspects of Modern Machine Learning Problems* **Fred Roosta-Khorasani** Crane Hodgkinson Tsuchida Eshragh	*MCMC Algorithms* **Michael Choi** Au Black Choi Tan	*Control of Stochastic Systems* **Hayriye Ayhan** Silva Gardner Gautam Drent	*Applied Probability in Financial Engineering* **Lingfei Li** Chen Gao Zhang Li
18:30 – 20:00	**Celebrating Women in APS** — At the Ship Inn (15-Minute Walk)			

Thursday July 4, 2019				
	Track 1	Track 2	Track 3	Track 4
	Rooms P1 and P2	Room P3	Room P4	Room P5
Session 4 *9:00–10:30*	*Tutorial #1* **Rhonda Righter & Kristy Gardner**	*Asymptotic Behaviour of Stochastic Systems* **Peter Braunsteins** McVinish Liu de Kemp Braunsteins	*Many-Server Queues* **David Stanford** Hermansson van Kreveld Toyoizumi Stanford	*Recent Advances in Stochastic Systems I* **Jing Dong** Arlotto Dai Maguluri Mendelson
10:30 – 11:00	**Morning Tea**			
11:00 – 12:00	**Keynote #3 (Ton Dieker)** (Rooms P1 and P2)			
12:00 – 13:00	**Lunch**			
Session 5 *13:00–14:30*	*Simulation, Risk, and Optimization* **Raghu Pasupathy** Henderson Deo Pasupathy Blanchet	*Stochastic Network Asymptotics* **Rami Atar** Budhiraja Lipshutz Ramanan Xie	*Queueing, Control, and Services* **Amy Ward** Gurvich Honnappa You Shi	*Recent Advances in Stochastic Systems II* **Alessandro Arlotto** Xu Dong Xu Sun
14:30 – 15:00	**Afternoon Tea**			
Session 6 *15:00–16:30*	*Open Problem Session* **Rami Atar & Harsha Honnappa**	*Large Scale Interacting Queueing Networks* **Ayalvadi Ganesh** Sankararaman Shneer Ziedins Ganesh	*Stochastic Models and Matching Queues* **Pascal Moyal** Ward Dorsman Weiss Moyal	*Approximations of Stochastic Processes* **Giang Nguyen & Oscar Peralta** Jansen Mandjes Nguyen Peralta
16:45 – 17:15	**APS AGM** (Rooms P1 and P2)			
18:00 – 21:00	**Conference Dinner** — At the Greek Club (5-Minute Walk)			

Thursday July 4, 2019				
	Track 5	Track 6	Track 7	Track 8
	Room M1	Room M2	Rooms M5 and M6	Rooms M7 and M8
Session 4 *9:00–10:30*	*New Developments in Reinforcement Learning* **Mengzhou Liu** Min Wang Chen Chen	*Diffusions and Related Stochastic Processes* **Jiyeon Lee** Jain Costantini Ethier Lee	*Optimal Stopping and Control* **David Goldberg** Russo Chen Goldberg Grangereau	*Healthcare I* **Zeynep Akşin Karaesmen** Yom-Tov Chen Örmeci James
10:30 – 11:00	**Morning Tea**			
11:00 – 12:00	**Keynote #3 (Ton Dieker)** (Rooms P1 and P2)			
12:00 – 13:00	**Lunch**			
Session 5 *13:00–14:30*	*Probabilistic Methods for Inference and Learning* **Mohsen Bayati** Goh Gur Russo Bayati	*Doubly Stochastic and State Dependent Queueing Models* **Andrew Daw** Gao Sun Badian-Pessot Daw	*Stochastic Optimal Control in Finance* **Duy-Minh Dang** Van Staden Lu Zhang Dang	*Healthcare II* **Ad Ridder** Liu Fackrell van Dijk Ridder
14:30 – 15:00	**Afternoon Tea**			
Session 6 *15:00–16:30*	*Learning in Sequential Decision Problems* **Dan Russo** Bayati Xu Shah Keppo	*Diffusions and Simulation* **Vlad Margarint** García-Pareja Shkolnik Margarint	*Monte Carlo Methods* **Pierre L'Ecuyer** Akin Bourgey L'Ecuyer	*Applications in Economics* **Nick Arnosti** Zhou Niknami Muir Arnosti
16:45 – 17:15	**APS AGM** (Rooms P1 and P2)			
18:00 – 21:00	**Conference Dinner** — At the Greek Club (5-Minute Walk)			

Friday July 5, 2019				
	Track 1	Track 2	Track 3	Track 4
	Rooms P1 and P2	Room P3	Room P4	Room P5
9:00 – 10:00	**Keynote #4 (Charles Bordenave)** (Rooms P1 and P2)			
10:00 – 10:30	**Morning Tea**			
Session 7 *10:30–12:00*	*Tutorial #2* **Devavrat Shah**	*Resource Allocation for Emerging Applications* **Siva Theja Maguluri** Hurtado-Lange Qian Zubeldia Gopalakrishnan	*Approximations and Controls for Queues* **Yunan Liu** Ward Liu Liu Sun	*Branching Random Walks* **Sophie Hautphenne & Peter Braunsteins** Zucca Bertacchi Hautphenne Braunsteins
12:00 – 13:00	**Lunch**			
Session 8 *13:00–14:30*	*Simulation Optimization* **Shane Henderson** Russo Zhou Hong Henderson	*Resource Allocation Models* **Alexander Stolyar** Shneer Zhong Xu Stolyar	*Statistical Inference in Queueing Systems* **Liron Ravner** Asanjarani Ravner Mandjes	*Recent Developments and Advances in Markov Decision Processes* **Sandjai Bhulai & Floske Spieksma** Kanavetas Bhulai He Spieksma
14:30 – 15:00	**Afternoon Tea**			
Session 9 *15:00–16:30*	*Convex Optimization in Applied Probability* **Chaithanya Bandi** Qian Maguluri Loke Proskynitopoulos	*Asymptotics and Algorithms on Random Graphs* **Mariana Olvera-Cravioto** Cao Fraiman Banerjee Olvera-Cravioto	*Queueing and Scheduling* **John Hasenbein** Adan Righter Hasenbein Weiss	*Jump Markov Processes* **Eugene Feinberg** Katehakis Spieksma Feinberg
16:30 – 17:00	**Cake and Farewell**			

Friday July 5, 2019				
	Track 5	Track 6	Track 7	Track 8
	Room M1	Room M2	Rooms M5 and M6	Rooms M7 and M8
9:00 – 10:00	**Keynote #4 (Charles Bordenave)** (Rooms P1 and P2)			
10:00 – 10:30	**Morning Tea**			
Session 7 *10:30–12:00*	*Information and Learning in Stochastic Systems* **Yuan Zhong** Xu Arlotto Russo Kamble	*Statistics and Stochastic Processes* **Sarat Moka** Ma Lewis Tafakori Moka	*Matching and Compensation* **Jussi Keppo** Loertscher Bhat Lim Lo	*Ecology/Block Chain* **Rhys Bowden** Warne Nitschke Wurm Bowden
12:00 – 13:00	**Lunch**			
Session 8 *13:00–14:30*	*Learning, Optimization, and Applied Probability* **Harsha Honnappa** Juneja Maguluri Zheng Pasupathy	*Large-Scale Inference and High-Dimensional Statistics* **Wen Sun** Radchenko Choi Xia Cope	*Storage Processes and Markov Chains* **Jaron Sanders** Jacobovic Deiana Ganguly Sanders	*Stochastic Modeling and Control in Healthcare and Service Operations* **Pengyi Shi** Peng Liu Dong Shi
14:30 – 15:00	**Afternoon Tea**			
Session 9 *15:00–16:30*	*Learning in Service Systems* **Stella Kapodistria** Peeters Saberi Kapodistria	*Random Walks/Dynamic Markets* **Ruiting Zuo** Abramov Kobayashi Muir Zuo	*Analysis of Computing and Service Systems* **Sherwin Doroudi** Abdul Jaleel Javier Gardner Doroudi	*Stochastic Modeling for Service Systems* **Nilay Tanik Argon & Lerzan Örmeci** Timmerman Ouyang Sun Örmeci
16:30 – 17:00	**Cake and Farewell**			

5

Plenaries

All plenaries take place in **Rooms P1 and P2**.

Wednesday 9:00 – 10:15 Sidney Resnick

Why Model the Growth of Networks?

Sidney Resnick (Marcel Neuts Lecturer), Cornell University, USA, sir1@cornell.edu

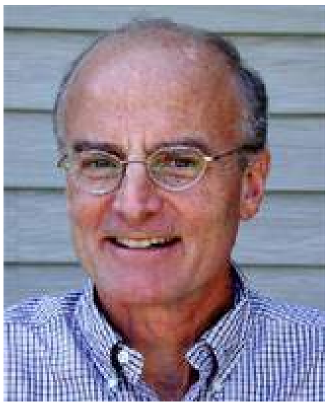

Social network modeling provides plenty of data but realistic models for network growth must be simple if mathematical results are expected. We have used preferential attachment (PA) models with a small number of parameters in an attempt to strike a balance between the mathematics and the statistical fitting. The PA models struggle to match the data but provide a context in which to test methods and analyze estimation techniques. Numerical summaries of network characteristics are often estimated using methods imported from classical statistics without real justification. For example, the Hill estimator coupled with a minimum distance threshold selection technique are commonly used. We discuss some attempts to justify and understand these estimation methods in the context of PA models. Without a model and its properties, there is no way to understand the limitation of estimation methods.

Wednesday 15:15 – 16:15 Nelly Litvak

Centrality in Large Random Networks

Nelly Litvak, University of Twente and Eindhoven University of Technology, the Netherlands, n.litvak@utwente.nl

A centrality is an algorithm that assigns an importance score to each node in a large network, such as an on-line social network or the World Wide Web. Most prominent example is the PageRank algorithm designed by Google to rank web pages. Mathematically, PageRank is a stationary distribution of a simple random walk with restart on a directed graph. A natural yet largely open question is: how the underlying graph structure affects the distribution of centrality scores? We will address this question by providing an overview of formal analysis for the PageRank distribution. One approach involves formulating the problem in terms of a stochastic fixed-point equation, which enables the analysis of PageRank on trees and on random graphs that can be coupled with a tree. Such random graphs include, e.g., the influential configuration model, where degrees of the nodes are fixed but connections are made purely at random. Another approach exploits the recent notion of local weak convergence of sparse graph sequences, and builds on the right balance between local and global properties of PageRank. We will discuss how this analysis helps to explain empirically observed phenomena and its potential extension to a large variety of more realistic network models.

The talk is based on joint work with M. Olvera-Cravioto, Y. Volkovich, N. Chen, R. van der Hofstad, A. Garavaglia.

Thursday 11:00 – 12:00 Ton Dieker

Towards a Next-Generation Methodology for Stochastic Network Analysis

Ton Dieker, Columbia University, USA, dieker@columbia.edu

We introduce a modeling methodology for stochastic networks that can be used to generate transient distributions of key performance metrics over time. We present a proof of concept that is lightning fast, remarkably accurate, and widely applicable: for example, it handles multi-server stations, time-varying arrivals, general service distributions, periods of significant overload, and re-entrant flow. Our approach circumvents the state space explosion required by a complete Markov representation by viewing the requisite information from a non-Markov probabilistic perspective. We will describe ideas for future enhancements that go far beyond our proof of concept.

Joint work with Steve Hackman (Georgia Tech).

Friday 9:00 – 10:00 Charles Bordenave

Non-backtracking Spectrum of Random Matrices

Charles Bordenave (IMS Medallion Lecturer), Université de Toulouse, France, charles.bordenave@univ-amu.fr

A fruitful line of thought in the study of discrete combinatorial structures, such as graphs, is to look for natural matrices or operators whose spectrum will contain valuable and accessible information on the underlying combinatorial structure. The non-backtracking matrix of a graph is a matrix acting on pairs of vertices sharing an edge. The entries of powers of the non-backtracking matrix count non-backtracking paths along the edges of the graph, that is, paths which do not visit twice the same edge successively. A non-backtracking path may be interpreted as a discrete geodesic. This matrix has been introduced by Hashimoto in 1988 in the context of the Ihara zeta function of a graph. In recent years, due to its strong geometric flavor, this non-backtracking matrix has been promoted as a powerful tool to analyze the subtle interplay between the geometry of graph and its spectrum. It has found a wide range of applications notably in spectral algorithms. We will focus the talk on the study of the extremal eigenvalues of non-backtracking matrices of classical random graph ensembles: uniform regular graphs, Erdős-Rényi random graphs, stochastic block models and random n-lifts of a base graph. We will notably explain how these results can be used in community detection problems and in the theory of expander graphs, classical and quantum.

This is based on joint works with Benoit Collins, Marc Lelarge and Laurent Massoulié.

6

Tutorials and Open Problem Session

All tutorials and the open problem session take place in **Rooms P1 and P2**.

Thursday 9:00 – 10:30 Rhonda Righter and Kristy Gardner

Product (Re)forms

Rhonda Righter, University of California, Berkeley, USA, rrighter@ieor.berkeley.edu
Kristy Gardner, Amherst College, USA, kgardner@amherst.edu

Recent years have seen a surge of interest in two distinct, but related queueing models featuring multiple classes of customers that differ in the set of servers at which they can be served. In one view of these systems, each arriving job joins the queue at all servers at which it can be served. In the "cancellation on start" (CoS) model, a job's extra copies are cancelled as soon as its first copy enters service, whereas in the "cancellation on completion" (CoC) model, a job's extra copies are cancelled as soon as its first copy completes service. Both models can also be viewed as having a single FCFS queue with either collaborative (CoC) or noncollaborative (CoS) service. The two systems are quite closely related mathematically, and surprisingly, under appropriate exponentiality assumptions, both yield a product-form stationary distribution. In this tutorial, we will discuss several different state-space descriptors for the two models, the product-form stationary distributions that result, and performance metrics that can be derived from the stationary distributions. We will also present new results comparing the performance of the two models, and we will discuss generalizations and related models, open questions, and directions for future research.

Friday 10:30 – 12:00 Devavrat Shah

Approximately Reversible Stochastic Processing Networks

Devavrat Shah, Massachusetts Institute of Technology, USA, devavrat@mit.edu

The property of (quasi-)reversibility of Markov chains have led to elegant characterization of steady-state distribution for complex queueing networks, e.g. celebrated Jackson networks, BCMP (Baskett, Chandi, Muntz, Palacois) and Kelly theorem. In a nutshell, despite the complicated interaction, in the steady-state, the queues in such networks exhibit independence and subsequently lead to explicit calculations of distributional properties of the queuing network that may seem impossible at the outset. The model of stochastic processing network (cf. Harrison (2000)) captures variety of dynamic resource allocation problems including the flow-level networks used for modeling bandwidth sharing in the Internet, switched networks (cf. Shah, Wischik (2006)) for modeling packet scheduling in the Internet router and wireless medium access, and hybrid flow-packet networks for modeling job-and-packet level scheduling in data centers. Unlike before, an appropriate resource allocation or scheduling policy is required in such networks to achieve good performance. Given the complexity, asymptotic analytic approaches such as fluid model or Lyapunov-Foster criteria to establish positive-recurrence and heavy traffic or diffusion approximation to characterize the scaled steady-state distribution became method of choice. A remarkable progress has been made along these lines over the past few decades, but there is a need for much more to match the explicit calculations in the context of reversible networks. In this tutorial, we will present an alternative to this approach that leads to non-asymptotic, explicit characterization of steady-state distribution akin BCMP / Kelly theorems. This involves (a) identifying a “relaxation” of the given stochastic processing network in terms of an appropriate (quasi-)reversible queueing network, and (b) finding a resource allocation or scheduling policy of interest that “emulates” the “relaxed” networks within “small error”. The proof is in the puddling – we will present three examples of this program: (i) distributed scheduling in wireless network, (ii) scheduling in switched networks, and (iii) flow-packet scheduling in a data center. The notion of “baseline performance” (cf. Harrison, Mandayam, Shah, Yang (2014)) will naturally emerges as a consequence of this program. We will discuss open questions pertaining multi-hop networks and computation complexity.

Thursday 15:00 – 16:30 Rami Atar and Harsha Honnappa

Open Problem Session

Rami Atar, Technion, Israel, rami@technion.ac.il
Harsha Honnappa, Purdue University, USA, honnappa@purdue.edu

The INFORMS APS 2019 conference introduces an open problem session, with the aim of providing an informal forum to present and discuss open problems on any subject in applied probability. The goal of this session is to stimulate ideas of new research directions by considering the gaps highlighted by the set of open problems. The structured part of this session consists of a series of speakers giving talks on open problems. The following speakers have been scheduled to present a talk:

Amarjit Budhiraja
University of North Carolina at Chapel Hill, USA, amarjit@unc.edu
Rare Event Asymptotics for JIQ and Related Systems

Eugene Feinberg
Stony Brook University, USA, eugene.feinberg@stonybrook.edu
Equivalence or Nonequivalence of Two Conditions for the Existence of Deterministic Average-Cost Optimal Policies for Infinite-State MDPs

Jerzy Filar
The University of Queensland, Australia, j.filar@uq.edu.au
Taxonomy of Uncertainties in Mathematical Models of the Environment

David Goldberg
Cornell University, USA, dag369@cornell.edu
Applied Probability and the ADP + ML Craziness

Raga Gopalakrishnan
Cornell University, USA, rg584@cornell.edu
Modeling, Analysis and Design of Service Systems with Strategic Agents

Sandeep Juneja
Tata Institute of Fundamental Research, India, juneja@tifr.res.in
Can an Algorithm that Matches the Error Probability Lower Bound be Designed for the Best Arm Multi-Armed Bandit Problem in the Fixed Budget Setting?

Amy Ward
The University of Chicago, USA, amy.ward@chicagobooth.edu
Performance Measure Approximation for Balking and Reneging Queues

Rami Atar
Technion, Israel, rami@technion.ac.il
A Diffusion Control Problem for Join-the-Shortest-Estimated-Queue

7

Sessions

Wednesday 10:45 – 12:15

Session 1.1 in Rooms P1 and P2 – Stochastic Simulation Optimization
Chair: Eunhye Song

1. *Rapid Search with Gaussian Markov Improvement Algorithm*
   Song, E.; Semelhago, M.; Nelson, B.; Wächter, A.
2. *An Adaptive Sequential Sample Average Approximation Solver for Two-stage Stochastic Programs*
   Pasupathy, R.; Song, Y.
3. *Green Simulation Optimization Using Likelihood Ratio Estimators*
   Eckman, D.; Feng, B.
4. *Parallel Adaptive Survivor Selection for VERY Large-Scale Simulation Optimization*
   Nelson, B.; Pei, L.; Hunter, S.

Session 1.2 in Room P3 – Stochastic Networks and Scaling Limits
Chair: Amarjit Budhiraja

1. *The divergence between Markovian and non-Markovian queueing models*
   Atar, R.; Budhiraja, A.; Dupuis, P.; Wu, R.
2. *Convergence rates to stationarity for reflecting Brownian motions*
   Banerjee, S.; Budhiraja, A.
3. *Asymptotic Behavior of a Critical Fluid Model for a Multiclass Processor Sharing Queue via Relative Entropy*
   Mulvany, J. A.; Puha, A. L.; Williams, R. J.
4. *Reflected Diffusions and (Bio)Chemical Reaction Networks*
   Anderson, D. F.; Higham, D. J.; Leite, S. C.; Williams, R. J.

Session 1.3 in Room P4 – Job Scheduling and Load Balancing in Queueing Systems
Chair: Onno Boxma & Jan-Pieter Dorsman

1. *Job Allocation in Large-Scale Service Systems with Affinity Relations*
   Cardinaels, E.; Borst, S.; van Leeuwaarden, J.
2. *Replicate to the shortest queues*
   Atar, R.; Keslassy, I.; Mendelson, G.
3. *Achievable stability in redundancy scheduling*
   Raaijmakers, Y.; Borst, S. C.; Boxma, O. J.
4. *Towards optimality in parallel job scheduling*
   Berg, B.; Dorsman, J. L.; Harchol-Balter, M.

Session 1.4 in Room P5 – Stochastic Systems
Chair: Mark Squillante & David Yao

1. *Why is Kemeny's constant a constant?*
   Bini, D.; Hunter, J.; Latouche, G.; Meini, B.; Taylor, P.
2. *Unique Ergodicity of Diffusion Limits of Many-Server Queues with Reneging*
   Agarwal, P.; Ramanan, K.
3. *Exponential Holding Cost in a Single-Server Queue: Interchange of Limits and Asymptotic Optimality of Reflection Control*
   Yang, J.; Yao, D. D.; Ye, H. Q.
4. *Fatou's lemma in its classic form for varying measures with applications to MDPs*
   Feinberg, E. A.; Kasyanov, P. O.; Liang, Y.

Session 1.5 in Room M1 – Statistical Methods for Learning and Decision-Making
Chair: Nan Ye

1. *Comparison of machine learning algorithms in classification of grazing behaviour in sheep*
   Hu, S.; Li, Y.; Ingham, A.; Hurley, G. B.; Gonzalez, E. G.; Wang, Y.-G.
2. *Efficient Newton-type Methods for Large-scale Invex Optimization*
   Liu, Y.; Roosta, F.
3. *Chaotic time series regression modeling using phase space reconstruction and deep neural network*
   Wang, Y. G.; Wu, J.
4. *Maximum Entropy Approaches for Inverse Reinforcement Learning*
   Snoswell, A. J.; Sing, S. P. N.; Ye, N.

Session 1.6 in Room M2 – Rare Events and Distributionally Robust Optimization
Chair: Chang-Han Rhee

1. *Confidence Regions for Optimal Transport Distributionally Robust Optimization*
   Blanchet, J.; Murthy, K.; Si, N.
2. *Safety evaluation of black-box prediction models via rare-event simulation*
   Huang, Z.; Lam, H.; Zhao, D.
3. *Sample path large deviation Levy processes and random walks with regularly varying increments in multiple dimensions*
   Su, Z.; Rhee, C.-H.
4. *Catastrophe principle for correlated stochastic processes*
   Rhee, C.-H.; Mazhba, M.; Chen, B.; Blanchet, J.; Zwart, B.

Session 1.7 in Rooms M5 and M6 – Control and Simulation of Stochastic Systems
Chair: Jiheng Zhang

1. *Perfect Sampling for Queues with Autoregressive Arrivals*
   Chen, X.
2. *Dynamic Routing in a Many-Server System*
   Cao, P.; Zhong, Z.; Huang, J.
3. *Dynamic Scheduling of Multiclass Many-server Queues with Abandonment: the Generalized $c\mu/h$ Rule*
   Long, Z.; Shimkin, N.; Zhang, H.; Zhang, J.
4. *An Asymptotically Optimal Index Policy for the Finite-Horizon Restless Bandit*
   Hu, W.; Frazier, P.; Zhang, X.

Session 1.8 in Rooms M7 and M8 – Stochastic Operations Research for Power Grids Applications
Chair: Alessandro Zocca

1. *Learning Graph Parameters from Linear Measurements: Fundamental Trade-offs and Application to Electric Grids*
   Li, T.; Werner, L.; Low, S. H.
2. *Optimal Switching Control of Continuous-Time Stochastic Systems: Theory and Applications*
   Martyr, R.
3. *No need to rush: dealing with deadlines in EV charging*
   Ferragut, A.; Zeballos, M.; Paganini, F.
4. *Less is More: Failure Localization in Power Systems*
   Zocca, A.; Guo, L.; Liang, C.; Low, S. H.; Wierman, A.

Wednesday 13:15 – 14:45

Session 2.1 in Rooms P1 and P2 – Model Uncertainty and Simulation Optimization
Chair: Enlu Zhou

1. *Sequential risk set inference for simulation optimization under input uncertainty*
   Song, E.
2. *Computationally Efficient Quantification of Simulation Input Uncertainty*
   Lam, H.; Qian, H.
3. *On the convergence rates of simulation with covariates*
   Li, C.; Gao, S.
4. *Online Quantification of Input Uncertainty for Parametric Models*
   Zhou, E.; Liu, T.

Session 2.2 in Room P3 – Queueing Networks and Markov Chains
Chair: Jim Dai

1. *State Space Collapse for Multi-class Queueing Networks with SBP Service Policies*
   Cao, C.; Dai, J. G.; Zhang, X. Y.
2. *Accelerating simulated annealing via replica exchange*
   Choi, M.
3. *Steady-state heavy traffic limits for multiclass queueing networks with SBP service policies*
   Braverman, A.; Dai, J. G.; Miyazawa, M.

Session 2.3 in Room P4 – Strategic Behavior in Queues
Chair: Moshe Haviv

1. *Coffee Shop Operations with Mobile Ordering*
   Kang, K.; Doroudi, S.; Delasay, M.
2. *Regulating Service Length Demand in a Single Server Queue*
   Oz, B.
3. *Equilibrium Threshold Strategy for An $M/M/1$ Feedback Queue*
   Fackrell, M.; Taylor, P.; Wang, J.
4. *On the (sub)optimality of the $c\mu$ rule when customers are strategic*
   Haviv, M.; Oz, B.

Session 2.4 in Room P5 – Asymptotic Analysis in Stochastic Systems
Chair: Guodong Pang

1. *Limit theorems for record indicators in threshold Nevzorov schemes*
   He, P.; Borovkov, K.
2. *Online Optimal Pricing And Capacity Sizing For A $G/G/1$ Queue With Demand Learning*
   Liu, Y.; Chen, X.
3. *Heavy-Traffic Limits for Stationary Network Flows*
   Whitt, W.; You, W.
4. *Nonstandard critical behavior of queue-based CSMA algorithms on the complete interference graph*
   Castiel, E.; Borst, S.; Miclo, L.; Simatos, F.; Whiting, P.

Session 2.5 in Room M1 – Learning and Bandits
Chair: Sandeep Juneja

1. *On The Futility of Dynamics In Robust Mechanism Design*
   Balseiro, S.; Kim, A.; Russo, D.
2. *Iterative Collaborative Filtering for Sparse Noisy Tensor Estimation*
   Shah, D.; Yu, C. L.
3. *Learning to detect an oddball target*
   Sundaresan, R.
4. *Best arm selection, and regret minimization, for general distributions, in the bandit framework*
   Agarwal, S.; Juneja, S.

Session 2.6 in Room M2 – Approximating Rare Event Probabilities
Chair: Offer Kella

1. *Counting Candy Crush Configurations*
   Hamilton, A.; Nguyen, G.; Roughan, M.
2. *Sampling Conditional on a Rare Event*
   Botev, Z.
3. *Rare Event Simulation for Steady-State Probabilities via Recurrency Cycles*
   Bisewski, K.; Crommelin, D.; Mandjes, M.
4. *Minimization of a stochastic convex function of a random variable subject to certain stochastic constraints*
   Kella, O.; Jacobovic, R.

Session 2.7 in Rooms M5 and M6 – Asymptotics and Control
Chair: Asaf Cohen

1. *Social Learning with Bandits*
   Sankararaman, A.; Ganesh, A.; Shakkottai, S.
2. *Enhanced balancing of bias-variance tradeoff in stochastic simulation*
   Lam, H.; Zhang, X.; Zhang, X.
3. *Sensitivity analysis of reflected diffusions*
   Lipshutz, D.; Ramanan, K.
4. *On singular control problems, the time-stretching method, and the weak-M1 topology*
   Cohen, A.

Session 2.8 in Rooms M7 and M8 – Revenue, Service, and Inventory Management
Chair: Qiong Wang & Marty Reiman

1. *Size Matters, So Does Duration: The Interplay between Offer Size and Offer Deadline*
   Hu, Z.; Tang, W.
2. *Asymptotic optimality of a local scheduling policy for a queueing system with customer feedback*
   Huang, J.
3. *A Re-solving Heuristic with Uniformly Bounded Loss for Network Revenue Management*
   Bumpensanti, P.; Wang, H.

Wednesday 16:30 – 18:00

Session 3.1 in Rooms P1 and P2 – Efficiency in Stochastic Optimization and Simulation Experiments
Chair: Henry Lam

1. *Efficient estimation of sensitivity of tail risk measures with applications in risk averse stochastic optimization*
   Deo, A.; Murthy, K.
2. *Space-filling design for non-linear models*
   Rhee, C.-H.; Zhou, E.; Qiu, P.
3. *Rare Event Estimation for Elliptical Random Vectors*
   Chen, D.; Pasupathy, R.; Honnappa, H.
4. *Bounding Optimality Gaps via Bagging*
   Lam, H.; Qian, H.

Session 3.2 in Room P3 – Dynamic Investment and Stochastic Networks
Chair: Agostino Capponi

1. *Dynamic Investment and Financing Decisions with Internal and External Liquidity Management*
   Chen, N.; Ji, J.; Tian, Y.
2. *A Dynamic Network Model of Interbank Lending – Systemic Risk and Liquidity Provisioning*
   Capponi, A.; Sun, X.; Yao, D. D.
3. *Disruption and Rerouting in Supply Chain Networks*
   Birge, J. R.; Capponi, A.; Chen, P.-C.
4. *A concentration phenomenon for system failure times under simultaneous shocks*
   Barrera, J.; Lagos, G.

Session 3.3 in Room P4 – Level-Dependent QBD Processes
Chair: Yonit Barron

1. *Drift Conditions for the Ergodicity of a Class of Level-Dependent $M/G/1$-type Processes*
   Cordeiro, J.; Kharoufeh, J.
2. *Asymptotic Analysis of QBDs with Time-Varying Periodic Rates with Examples*
   Margolius, B.
3. *Extremal Queues Given First Two Moments*
   Chen, Y.; Whitt, W.
4. *A threshold policy in a Markov-modulated production system with server vacation: The case of continuous and batch supplies*
   Barron, Y.

Session 3.4 in Room P5 – Markov Chains
Chair: Ivo Adan

1. *Overview of first passage times calculations for double thresholded processes*
   Kapodistria, S.
2. *Algebraic Classification of Markov Chains for Assessing Quasi-Birth–Death Process Tractability*
   Wickeham, A.; Doroudi, S.
3. *Stability condition for a two class discrete-time queueing model with bounded variable-length service times*
   Dendievel, S.; Bruneel, H.
4. *Optimizing Activity Sequences by Minimizing Rework Effects*
   Adan, I.; Wilschut, T.

Session 3.5 in Room M1 – Statistical and Computational Aspects of Modern Machine Learning Problems
Chair: Fred Roosta-Khorasani

1. *Communication-Efficient Distributed Optimization Methods for Large-Scale Machine Learning*
   Crane, R.; Roosta, F.
2. *Implicit Langevin algorithms for sampling from log-concave densities*
   Hodgkinson, L.; Salomone, R.; Roosta, F.
3. *Probabilistic Symmetries in Neural Networks*
   Tsuchida, R.; Roosta, F.; Gallagher, M.
4. *A new fast algorithm to approximate the leverage scores of big time series data: theory and application*
   Eshragh, A.; Roosta, F.; Mahoney, M.; Nazari, A.

Session 3.6 in Room M2 – MCMC Algorithms
Chair: Michael Choi

1. *Constrained Hamiltonian Monte Carlo for PDE Inverse Problems*
   Au, K. X.; Graham, M. M.; Thiery, A. H.
2. *Importance sampling for partially observed Markov chain models*
   Black, A. J.
3. *Hitting time, mixing time and geometric interpretations of Metropolis-Hastings reversiblizations*
   Choi, M.; Huang, L.
4. *Equity Valuation and Falling Cost of Data Analytics*
   Keppo, J; Tan, H. M.

Session 3.7 in Rooms M5 and M6 – Control of Stochastic Systems
Chair: Hayriye Ayhan

1. *Scheduling Same-day Deliveries Using Truck-to-Drone Re-supply*
   Silva, D. F.; Smith, A.
2. *A Little Redundancy Goes A Long Way: Convexity in Redundancy Systems*
   Gardner, K.; Hyytiä, E.; Righter, R.
3. *Analysis of Batch Dispatch Policies for Order Fulfillment*
   Gautam, N.
4. *Realtime Integrated Learning and Decision Making for Deteriorating Systems*
   Drent, C.; Drent, M.; Arts, J; Kapodistria, S.

Session 3.8 in Rooms M7 and M8 – Applied Probability in Financial Engineering
Chair: Lingfei Li

1. *A multi-factor regime switching model for inter-trade durations in the limit order market*
   Chen, X.; Li, Z.; Xing, H.
2. *Langevin Algorithms and Momentum-Based Acceleration for Non-convex Stochastic Optimisation*
   Gao, X.; Gurbuzbalaban, M.; Zhu, L.
3. *Analysis of Markov Chain Approximation for Diffusion Models with Non-Smooth Coefficients*
   Li, L.; Zhang, G.
4. *A General Method for Valuation of Drawdown Risk under Markovian Models*
   Zhang, G.; Li, L.

Thursday 9:00–10:30

Session 4.2 in Room P3 – Asymptotic Behaviour of Stochastic Systems
Chair: Peter Braunsteins

1. *Fast approximate simulation of long-range finite spin systems*
   McVinish, R.; Hodgkinson, L.
2. *On the asymptotic behaviour of the number of renewals via translated Poisson*
   Xia, A.; Liu, Q.
3. *Performance of the smallest-variance-first rule in appointment sequencing*
   de Kemp, M. A.; Mandjes, M.; Olver, N.
4. *Local limit theorems for occupancy models*
   Barbour, A.; Braunsteins, P.; Ross, N.

Session 4.3 in Room P4 – Many-Server Queues
Chair: David Stanford

1. *Monotonicities in systems of parallel Processor Sharing queues*
   Hermansson, N.; Ziedins, I.
2. *A symmetric Erlang loss queue with breakdowns*
   van Kreveld, L.; Boxma, O. J.; Dorsman, J. L.; Mandjes, M.
3. *The Battle between Infinite Server Queues and Heavy-tailed Arrivals*
   Toyoizumi, H.
4. *The average waiting time for both classes in a delayed accumulating priority queue*
   Stanford, D. A.; Bilodeau, B.

Session 4.4 in Room P5 – Recent Advances in Stochastic Systems I
Chair: Jing Dong

1. *Uniformly Bounded Regret in the Multi-Secretary Problem*
   Arlotto, A.; Gurvich, I.
2. *High-order accuracy steady-state diffusion approximations*
   Dai, J.
3. *Heavy-Traffic Analysis of Queueing Systems with no Complete Resource Pooling*
   Hurtado-Lange, D.; Mou, S.; Maguluri, S. T.
4. *Sub-diffusive load-balancing in time-varying queueing systems*
   Atar, R.; Keslassy, I.; Mendelson, G.

Session 4.5 in Room M1 – New Developments in Reinforcement Learning
Chair: Mengzhou Liu

1. *Thompson Sampling with Information Relaxation Penalties*
   Min, S.; Maglaras, C.; Moallemi, C. C.
2. *Continuous-Time Mean-Variance Portfolio Selection: A Reinforcement Learning Framework*
   Wang, H.; Zhou, X.
3. *Duality-Based Exploration in Reinforcement Learning and Its Application*
   Chen, N.; Ma, X.
4. *Dynamic Pricing of Relocating Resources in Large Networks*
   Balseiro, S. R.; Brown, D. B.; Chen, C.

Session 4.6 in Room M2 – Diffusions and Related Stochastic Processes
Chair: Jiyeon Lee

1. *Probabilistic Contraction Analysis of Iterated Random Operators*
   Gupta, A.; Jain, R.; Glynn, P.
2. *Reflecting diffusions in nonsmooth domains: Markov selection and uniqueness*
   Costantini, C.; Kurtz, T. G.
3. *The tilted flashing Brownian ratchet*
   Ethier, S. N.; Lee, J.
4. *How strong can the Parrondo effect be?*
   Ethier, S. N.; Lee, J.

Session 4.7 in Rooms M5 and M6 – Optimal Stopping and Control
Chair: David Goldberg

1. *A Finite Time Analysis of Temporal Difference Learning*
   Bhandari, J.; Russo, D.; Singal, R.
2. *A new approach to high-dimensional online decision-making*
   Chen, Y.; Goldberg, D. A.
3. *Beating the curse of dimensionality in options pricing and optimal stopping*
   Goldberg, D. A.; Chen, Y.
4. *McKean stochastic optimal control of an energy storage system to reduce demand variability*
   Grangereau, M.; Gobet, E.

Session 4.8 in Rooms M7 and M8 – Healthcare I
Chair: Zeynep Akşin Karaesmen

1. *Balancing Multi-Featured Workload in Service Systems*
   Carmeli, N.; Yom-Tov, G. B.; Mandelbaum, A.
2. *Sizing flexible resources in large-scale service systems*
   Chen, J.; Dong, J.
3. *Behavior of strategic patients in a health care system with public and private facilities*
   Örmeci, L.; Andritsos, D.; Canbolat, P.; Dimitrakopoulos, Y.; Sadeghzadeh, S.
4. *Modelling intensive care units using quasi-birth-death processes*
   James, S.; Bean, N.; Tuke, J.

Thursday 13:00–14:30

Session 5.1 in Rooms P1 and P2 – Simulation, Risk, and Optimization
Chair: Raghu Pasupathy

1. *An Envelope Procedure for Ranking and Selection Problems*
   Ma, S.; Henderson, S. G.
2. *Systemic Risk - A simulation perspective*
   Deo, A.; Juneja, S.
3. *Convergence Rates of Sampling Algorithms for Stochastic Optimization: The Effect of Curvature*
   Newton, D.; Pasupathy, R.; Yip, A. N. K.
4. *Optimal Transport Based Distributionally Robust Optimization: Structural Properties and Iterative Schemes*
   Blanchet, J.; Murthy, K.; Zhang, F.

Session 5.2 in Room P3 – Stochastic Network Asymptotics
Chair: Rami Atar

1. *JSQ, LDP, and the Golden Ratio*
   Budhiraja, A.; Friedlander, E.; Wu, R.
2. *Heavy traffic limits for load balancing algorithms using delayed information*
   Atar, R.; Lipshutz, D.
3. *Interacting Markov chains on Large Sparse Graphs*
   Ganguly, A.; Lacker, D.; Ramanan, K.; Wu, R.
4. *Logarithmic Regret in the Dynamic and Stochastic Knapsack Problem with Equal Rewards*
   Arlotto, A.; Xie, X.

Session 5.3 in Room P4 – Queueing, Control, and Services
Chair: Amy Ward

1. *A Moment Matching Algorithm for Dynamic Programs*
   Gurvich, I.; Zhang, B.-Z.
2. *Computational Algorithms for Optimal Design of Service Systems*
   Jaiswal, P.; Honnappa, H.; Rao, V. A.
3. *Time-Varying Robust-Queueing for Quantile Approximation*
   Whitt, W.; You, W.
4. *Timing it Right: Balancing Inpatient Congestion versus Readmission Risk at Discharge*
   Shi, P.; Helm, J.; Deglise-Hawkinson, J.; Pan, J.

Session 5.4 in Room P5 – Recent Advances in Stochastic Systems II
Chair: Alessandro Arlotto

1. *Efficient random graph matching via degree profiles*
   Ding, J.; Ma, Z.; Wu, Y.; Xu, J.
2. *The Power of Two in Queue Scheduling*
   Chen, Y.; Dong, J.
3. *Query Complexity of Bayesian Private Learning*
   Xu, K.
4. *Optimal Contract for Machine Repair and Maintenance*
   Tian, F.; Sun, P.; Duenyas, I.

Session 5.5 in Room M1 – Probabilistic Methods for Inference and Learning
Chair: Mohsen Bayati

1. *Beyond Binary Inference: Contrast-Specific Propensity Scores for Heterogeneous Treatment Effect Estimation*
   Han, S.; Goh, J.; Meng, F.; Rubin, D. B.
2. *Adaptive Sequential Experiments with Unknown Information Flows*
   Gur, Y.; Momeni, A.
3. *Global Convergence Guarantees for Policy Gradient Methods*
   Bhandari, J.; Russo, D.
4. *A General Tail Bound for Matrix Estimation*
   Hamidi, N.; Bayati, M.

Session 5.6 in Room M2 – Doubly Stochastic and State Dependent Queueing Models
Chair: Andrew Daw

1. *Affine Point Processes: Refinements to Large-Time Asymptotics*
   Gao, X.; Zhu, L.
2. *Many-Server Queues with Autoregressive Inputs*
   Sun, X.
3. *Optimal Control Policies for an $M/M/1$ Queue with a Removable Server and Dynamic Service Rates*
   Badian-Pessot, P.; Down, D.; Lewis, M.
4. *The Queue-Hawkes Process: Ephemeral Self-Excitement*
   Daw, A.; Pender, J.

Session 5.7 in Rooms M5 and M6 – Stochastic Optimal Control in Finance
Chair: Duy-Minh Dang

1. *Mean-Quadratic Variation (MQV) portfolio optimisation as an alternative to Time-consistent Mean-Variance (TCMV) optimisation*
   Van Staden, P. M.; Dang, D. M.; Forsyth, P. A.
2. *Numerical methods for Guaranteed Minimum Withdrawal Benefits (GMWBs) as a continuous impulse control problem*
   Lu, Y.; Dang, D. M.; Forsyth, P. A.
3. *An integration method for mean-variance portfolio optimisation under the jump-extended Heston model*
   Zhang, H.; Dang, D.
4. *Mean-risk portfolio optimisation*
   Dang, D. M.

Session 5.8 in Rooms M7 and M8 – Healthcare II
Chair: Ad Ridder

1. *Modelling of reporting behaviour in the FluTracking surveillance system*
   Liu, D.; Mitchell, L.; Carlson, S.; Ross, J. V.
2. *Increasing the Health Benefits of In-home Chronic Care by Optimal Service Timing*
   Fackrell, M.; Fu, J.; Taylor, P. G.; Tirdad, A.
3. *On congestion probabilities for ICU-SDU systems with overflow*
   van Dijk, N. M.; Schilstra, B. H.
4. *Minimizing bed occupancy variance by scheduling patients under uncertainty*
   Ridder, A.; van den Broek d'Obrenan, A.; Roubos, D.; Stougie, L.

Thursday 15:00–16:30

Session 6.2 in Room P3 – Large Scale Interacting Queueing Networks
Chair: Ayalvadi Ganesh

1. *Interference Queueing Networks on Grids*
   Sankararaman, A.; Baccelli, F.; Foss, S.
2. *Stability under utility-maximising allocations: from finite to infinite networks*
   Shneer, S.; Stolyar, A.
3. *Phase transitions in a generalized hard core model (k - casting)*
   Ramanan, K.; Ziedins, I.
4. *Large deviations for Cox processes and $Cox/G/\infty$ queues, with a biological application*
   Dean, J.; Ganesh, A.; Crane, E.

Session 6.3 in Room P4 – Stochastic Models and Matching Queues
Chair: Pascal Moyal

1. *Dynamic Matching for Real-Time Ridesharing*
   Ward, A. R.; Özkan, E.
2. *A token-based central queue with order-independent service rates*
   Ayesta, U.; Bodas, T.; Dorsman, J. L.; Verloop, I. M.
3. *FCFS Dynamic Matching Models*
   Weiss, G.
4. *Coupling from the past of stochastic matching models*
   Moyal, P.; Busic, A.; Mairesse, J.

Session 6.4 in Room P5 – Approximations of Stochastic Processes
Chair: Giang Nguyen & Oscar Peralta

1. *The Asymptotic Behavior of the Time-Varying Supermarket Model*
   Jansen, H. M.
2. *Simulation-based assessment of the stationary tail distribution of a stochastic differential equation*
   Bisewski, K.; Crommelin, D.; Mandjes, M.
3. *Rate of strong convergence of stochastic fluid processes to Markov-modulated Brownian motion*
   Nguyen, G.; Peralta, O.
4. *Convergence of a bivariate flip–flop process*
   Latouche, G.; Nguyen, G.; Peralta, O.

Session 6.5 in Room M1 – Learning in Sequential Decision Problems
Chair: Dan Russo

1. *Contextual Bandits with a Low-Rank Structure*
   Hamidi, N.; Bayati, M.; Gupta, K.
2. *Delay-Predictability tradeoffs in reaching a secret goal*
   Tsitsiklis, J. N.; Xu, K.
3. *AlphaGo Zero, Monte Carlo Tree Search and Self-Play: Towards Theoretical Foundations*
   Shah, D.; Xie, Q.; Xu, Z.
4. *Costly Learning Manipulation: Belief Distortion Through Information Dissemination*
   Keppo, J.; Kim, M.; Zhang, X.

Session 6.6 in Room M2 – Diffusions and Simulation
Chair: Vlad Margarint

1. *Exact Simulation of Coupled Wright–Fisher Diffusions*
   García-Pareja, C.; Hult, H.; Koski, T.
2. *Unbiased Simulation of Multivariate Jump-Diffusions*
   Chen, G.; Giesecke, K.; Shkolnik, A.
3. *Pathwise and probabilistic analysis in the context of SLE*
   Margarint, V.

Session 6.7 in Rooms M5 and M6 – Monte Carlo Methods
Chair: Pierre L'Ecuyer

1. *Control variates for censored Monte Carlo simulation*
   Akin, E.; Szechtman, R.
2. *Multilevel Monte-Carlo method and lower/upper bounds in Initial Margin computations*
   Bourgey, F.; De Marco, S.; Gobet, E.; Zhou, A.
3. *Quasi-Monte Carlo Density Estimation*
   Ben Abdellah, A.; L'Ecuyer, P.; Owen, A. B.; Puchhammer, F.

Session 6.8 in Rooms M7 and M8 – Applications in Economics
Chair: Nick Arnosti

1. *Misconducts in Organizations*
   Zhou, M.; Keppo, J.; Jokivuolle, E.
2. *Modelling Double Auctions with dynamic supply and Demand*
   Niknami, B.; Taylor, P. G.
3. *Optimal Market Thickness and Clearing*
   Loertscher, S.; Muir, E. V.; Taylor, P. G.
4. *A Continuum Model of Stable Matching with Finite Capacities*
   Arnosti, N.

Friday 10:30–12:00

Session 7.2 in Room P3 – Resource Allocation for Emerging Applications
Chair: Siva Theja Maguluri

1. *Transform Methods for Heavy-Traffic Analysis*
   Hurtado-Lange, D.; Maguluri, S. T.
2. *State Dependent Control of Closed Queueing Networks with Application to Ride-Hailing*
   Banerjee, S.; Kanoria, Y.; Qian, P.
3. *Delay and stability in distributed service systems with redundancy and random slowdowns*
   Gamarnik, D.; Tsitsiklis, J. N.; Zubeldia, M.
4. *On the design of service systems when servers are strategic*
   Gopalakrishnan, R.

Session 7.3 in Room P4 – Approximations and Controls for Queues
Chair: Yunan Liu

1. *A Fluid Limit for an Overloaded Multiclass Many-Server Queue with General Reneging Distribution*
   Puha, A.; Ward, A. R.
2. *Admission control for double-ended queueing systems in heavy traffic*
   Liu, X.; Weerasinghe, A.
3. *Staffing and Scheduling to Differentiate Service in Multiclass Time-Varying Service Systems*
   Liu, Y.; Sun, X.; Hovey, K.
4. *Optimal Repositioning of Urban Electric Vehicle Sharing Systems*
   Sun, X.

Session 7.4 in Room P5 – Branching Random Walks
Chair: Sophie Hautphenne & Peter Braunsteins

1. *The global critical value of branching random walks: weak but hard!*
   Zucca, F.
2. *The cardinality of extinction probability vectors of a branching random walk*
   Bertacchi, D.; Zucca, F.
3. *Extinction in lower Hessenberg branching processes with countably many types*
   Braunsteins, P.; Hautphenne, S.
4. *The probabilities of extinction in a branching random walk on a strip*
   Braunsteins, P.; Hautphenne, S.

Session 7.5 in Room M1 – Information and Learning in Stochastic Systems
Chair: Yuan Zhong

1. *Experimenting in Equilibrium*
   Wager, S.; Xu, K.
2. *Dynamic and stochastic knapsack problems: heuristics, asymptotic performance, and limit theorems*
   Arlotto, A.; Xie, X.
3. *Deep Exploration via Randomized Value Functions*
   Osband, I.; Van Roy, B.; Russo, D.; Wen, Z.
4. *Individual Fairness in Hindsight*
   Gupta, S.; Kamble, V.

Session 7.6 in Room M2 – Statistics and Stochastic Processes
Chair: Sarat Moka

1. *Modelling structural breaks in Autoregressive Time Series using parametric spectral discrimination tests*
   Grant, A.; Ma, L.; Sofronov, G.; Wishart, J. R.
2. *Algorithms for inferring parameters and hidden states of Markovian-regime-switching models with independent regimes for electricity prices*
   Lewis, A.; Bean, N.; Nguyen, G.
3. *Distance covariance for discretized stochastic processes*
   Dehling, H.; Matsui, M.; Mikosch, T.; Samorodnitsky, G.; Tafakori, L.
4. *Unbiased Estimation of Reciprocal of the Mean for non-negative Random Variable*
   Moka, S.; Kroese, D.; Juneja, S.

Session 7.7 in Rooms M5 and M6 – Matching and Compensation
Chair: Jussi Keppo

1. *Efficient adjustment dynamics*
   Loertscher, S.; Muir, E.
2. *Double Dipping of Two-Sided Platform Economy*
   Bhat, S.; Keppo, J.; Teo, C. P.
3. *Approximating the Gittins Index in a Bayesian bandit*
   Lim, A.
4. *Mechanisms for Incentive-Compatible Information-Sharing in Informal Supply Chains*
   de Zegher, J.; Lo, I.

Session 7.8 in Rooms M7 and M8 – Ecology/Block Chain
Chair: Rhys Bowden

1. *Applications of Bayesian inference, model selection, and experimental design in cell biology*
   Warne, D. J.; Baker, R. E.; Simpson, M. J.
2. *Stochastic Models of Multilevel Darwinian Populations*
   Nitschke, M. C.; Black, A.; Rainey, P.
3. *Polyp fiction: a stochastic fluid flow model for coral-algal symbiosis on the Great Barrier Reef*
   Wurm, M.; Bean, N.; Nguyen, G.
4. *Scaling and consensus in cryptocurrency blockchains*
   Bowden, R. A.; Krzesinski, A. E.; Taylor, P. G.

Friday 13:00–14:30

Session 8.1 in Rooms P1 and P2 – Simulation Optimization
Chair: Shane Henderson

1. *Simple Bayesian Algorithms for Best Arm Identification*
   Russo, D.
2. *Ranking and Selection under Input Uncertainty*
   Wu, D.; Zhou, E.
3. *Knockout-Tournament Procedures for Large-Scale Ranking and Selection in Parallel Computing Environments*
   Zhong, Y.; Hong, L.
4. *Fixed-confidence, fixed-tolerance guarantees for ranking-and-selection procedures*
   Eckman, D. J.; Henderson, S. G.

Session 8.2 in Room P3 – Resource Allocation Models
Chair: Alexander Stolyar

1. *Stability and moment bounds under utility-maximising service allocations*
   Shneer, S.; Stolyar, A.
2. *Improved queue-size scaling for input-queued switches via graph factorization*
   Xu, J.; Zhong, Y.
3. *Information and Memory in Dynamic Resource Allocation*
   Xu, K.; Zhong, Y.
4. *Discrete-time TASEP with holdback*
   Shneer, S.; Stolyar, A.

Session 8.3 in Room P4 – Statistical Inference in Queueing Systems
Chair: Liron Ravner

1. *A Survey of Parameter and State Estimation in Queues*
   Asanjarani, A.; Nazarathy, Y.; Taylor, P.
2. *Statistical estimation of the input to a queue by Poisson probing*
   Ravner, L.; Boxma, O.; Mandjes, M.
3. *Hypothesis testing in queueing*
   Mandjes, M.; Ravner, L.

Session 8.4 in Room P5 – Recent Developments and Advances in Markov Decision Processes
Chair: Sandjai Bhulai & Floske Spieksma

1. *Asymptotically optimal control for Markov Decision Processes (MDP) under side constraints*
   Burnetas, A.; Kanavetas, O.; Katehakis, M.
2. *Data-driven consumer debt collection via machine learning and approximate dynamic programming*
   van de Geer, R.; Wang, Q.; Bhulai, S.
3. *Optimal Policies for Stochastic Clearing Systems with Time-Dependent Delay Penalties*
   He, Q.-M.; Bookbinder, J. H.; Cai, Q.
4. *Bounds for threshold and switching curve optimal policies*
   Spieksma, F. M.

Session 8.5 in Room M1 – Learning, Optimization, and Applied Probability
Chair: Harsha Honnappa

1. *Discriminative Learning via Adaptive Questioning*
   Bassamboo, A.; Juneja, S.; Zeevi, A.
2. *Distributed Learning in Multi-Agent Systems: Optimization and Reinforcement Learning*
   Doan, T. T.; Maguluri, S. T.; Romberg, J.
3. *Stochastic Modeling and Decisions for Online-to-offline Supermarket Retails*
   Cao, J.; Shen, Z.; Zheng, Z.
4. *Stochastic Trust Region Methods ASTRO and ASTRO-DF: Convergence, Complexity, and Numerical Performance*
   Vasquez, D.; Pasupathy, R.; Shashaani, S.

Session 8.6 in Room M2 – Large-Scale Inference and High-Dimensional Statistics
Chair: Wen Sun

1. *High-dimensional regression with discrete optimization*
   Hazimeh, H.; Mazumder, R.; Radchenko, P.
2. *Community Detection via L-1 Fused Penalty*
   Choi, Y.; Tan, V.; Liu, Z.
3. *Privacy Preserving Integrative Regression Analysis of High-dimensional Heterogeneous Data*
   Xia, Y.
4. *Estimating the probability of emergence of a novel disease strain against an endemic background*
   Cope, R.; Meehan, M.

Session 8.7 in Rooms M5 and M6 – Storage Processes and Markov Chains
Chair: Jaron Sanders

1. *Steady-state optimization of an exhaustive Lévy storage process with intermittent output and random output rate*
   Jacobovic, R.; Kella, O.
2. *Sojourn time distribution in fluid queues*
   Deiana, E.; Latouche, G.; Remiche, M.-A.
3. *Limit Theorems for Markov Chains on Sparse Graphs*
   Ganguly, A.; Ramanan, K.
4. *Clustering in Block Markov Chains*
   Sanders, J.; Proutière, A.; Yun, S.-Y.

Session 8.8 in Rooms M7 and M8 – Stochastic Modeling and Control in Healthcare and Service Operations
Chair: Pengyi Shi

1. *Managing Time-varying Arrivals in Service Systems: A Dual Approach*
   Peng, X.
2. *To Pool or Not to Pool: Queueing Design for Large-Scale Service Systems with Customer Abandonment*
   Cao, P.; He, S.; Huang, J.; Liu, Y.
3. *Optimal Scheduling of Proactive Care with Customer Degradation*
   Hu, Y.; Chan, C.; Dong, J.
4. *Integrating New Diagnostic Tests Into Emergency Department Workflow*
   Helm, J.; Shi, P.; Heese, S.; Mitchell, A.

Friday 15:00–16:30

Session 9.1 in Rooms P1 and P2 – Convex Optimization in Applied Probability
Chair: Chaithanya Bandi

1. *Near Optimal Control of a Ride-Hailing Platform via Mirror Backpressure*
   Qian, P.; Kanoria, Y.
2. *Dynamic Pricing in Two Sided Queues*
   Varma, S.; Bumpensanti, P.; Wang, H.; Maguluri, S. T.
3. *Exploiting Hidden Convexity for Optimal Flow Control in Queueing Networks*
   Bandi, C.; Loke, G. G.
4. *Robust Queue Inference Engine*
   Bandi, C.; Proskynitopoulos, A.

Session 9.2 in Room P3 – Asymptotics and Algorithms on Random Graphs
Chair: Mariana Olvera-Cravioto

1. *Connectivity of a general class of inhomogeneous random digraphs*
   Cao, J.; Olvera-Cravioto, M.
2. *Recursive functions on conditional Galton–Watson trees*
   Broutin, N.; Devroye, L.; Fraiman, N.
3. *Non-parametric change point detection in growing networks*
   Banerjee, S.; Bhamidi, S.; Carmichael, I.
4. *PageRank under degree correlations*
   Olvera-Cravioto, M.

Session 9.3 in Room P4 – Queueing and Scheduling
Chair: John Hasenbein

1. *Newsvendor Equations for Production Networks*
   Adan, I.; Atan, Z.; Jansen, S.; de Kok, A.
2. *The Potentially Negative Effects of Cooperation in Service Systems*
   Chung, H.; Ahn, H.; Righter, R.
3. *A New Askew View of* c-μ
   Hasenbein, J.; Invernizzi, P.; Ripanti, S.
4. *Fluid models for queueing networks*
   Weiss, G.

Session 9.4 in Room P5 – Jump Markov Processes
Chair: Eugene Feinberg

1. *On Adaptive Control for Continuous-Time Markov Decision Processes*
   Cowan, W.; Ghosh, D.; Kanavetas, O.; Katehakis, M.; Papadimitriou, S.; Pirutinsky, D.
2. *Non-singular rate matrices in the successive lumping method for Quasi-Skipfree processes*
   Spieksma, F. M.; Ertiningsih, D.
3. *Kolmogorov's forward equation for jump Markov processes and its application to control problems*
   Feinberg, E. A.; Mandava, M.; Shiryaev, A. N.

Session 9.5 in Room M1 – Learning in Service Systems
Chair: Stella Kapodistria

1. *Continuous assortment optimization with logit choice probabilities*
   Peeters, Y.; den Boer, A.; Mandjes, M.
2. *Managing Ultra-Fast Delivery Service Using Personalized Promotion Strategy for Online*
   Saberi, Z.; Hussain, O.; Saberi, M.; Chang, E.
3. *Decision making under uncertainty*
   Kapodistria, S.

Session 9.6 in Room M2 – Random Walks/Dynamic Markets
Chair: Ruiting Zuo

1. *Conditions for recurrence and transience for one family of random walks*
   Abramov, V. M.
2. *Geometric properties of moment generating functions of increments for a three dimensional reflecting random walk*
   Kobayashi, M.
3. *Dynamic Market-Making: A Myersonian Approach*
   Akbarpour, A.; Muir, E. V.; Strack, P.
4. *Dynamic Contracting in Worker-Manager-Owner Relationship*
   Keppo, J.; Touzi, N.; Zuo, R.

Session 9.7 in Rooms M5 and M6 – Analysis of Computing and Service Systems
Chair: Sherwin Doroudi

1. *Scalable Load Balancing in the Presence of Servers Under Interference*
   Abdul Jaleel, J.; Doroudi, S.; Delasay, M.
2. *An Exact Analysis of a Class of Markovian Bitcoin Models*
   Javier, K.; Fralix, B.
3. *Dispatching in Heterogeneous Systems*
   Gardner, K.; Stephens, C.
4. *Dynamic Market Prices for Call Centers with Co-Sourcing Subject to Non-stationary Demand*
   Doroudi, S.; Mousavi, M.

Session 9.8 in Rooms M7 and M8 – Stochastic Modeling for Service Systems
Chair: Nilay Tanik Argon & Lerzan Örmeci

1. *Halfin-Whitt scaling of the Fixed-Cycle Traffic-Light queue*
   Boon, M. A. A.; Janssen, A. J. E. M.; van Leeuwaarden, J. S. H.; Timmerman, R. W.
2. *Assigning Priorities (or not) in Service Systems with Nonlinear Waiting Costs*
   Ouyang, H.; Argon, N.; Ziya, S.
3. *When to Triage in Service Systems with Hidden Customer Class Identities?*
   Sun, Z.; Argon, N. T.; Ziya, S.
4. *Strategic Customers in Systems with Batch Arrivals*
   Örmeci, L.; Bountali, O.; Burnetas, A.

8

Abstracts

Wednesday 10:45 – 12:15

Session 1.1 in Rooms P1 and P2 – Stochastic Simulation Optimization
Chair: Eunhye Song

Rapid Search with Gaussian Markov Improvement Algorithm

E. Song, Pennsylvania State University, United States, eus358@psu.edu
M. Semelhago, Northwestern University, United States,marksemelhago2019@u.northwestern.edu
B. Nelson, Northwestern University, United States, nelsonb@northwestern.edu
A. Wächter, Northwestern University, United States, waechter@iems.northwestern.edu

Gaussian Markov Improvement Algorithm (GMIA) is a simulation optimization procedure for a problem defined on a discrete feasible solution space. GMIA models the objective function values at the feasible solutions as a realization of a Gaussian Markov random field (GMRF) and uses it to compute the sampling criterion, complete expected improvement (CEI) relative to the current optimum, to sequentially select the next solutions to simulate. When the feasible solution space is large, computing the CEIs of all solutions becomes a bottleneck for GMIA, whereas most solutions' CEI values do not change much from iteration to iteration. In this talk, we introduce rapid-GMIA (r-GMIA) that alternates between rapid search and global search phases to balance computational efficiency and search effectiveness. At each global search iteration, a set of promising solutions is constructed. During the rapid search iterations, the exact CEIs of all solutions within the set are computed and the next simulation decision is made within the set. Per-iteration computational cost of rapid search phase is negligible compared to that of GMIA resulting a faster search progress for r-GMIA.

An Adaptive Sequential Sample Average Approximation Solver for Two-stage Stochastic Programs

R. Pasupathy, Purdue University, United States, pasupath@purdue.edu
Y. Song, Clemson University, United States, yongjis@clemson.edu

We present adaptive sequential SAA algorithms to solve large-scale two-stage stochastic linear programs. The iterative framework is organized as follows: during each outer iteration, a sample-path problem is implicitly generated using a sample of observations or "scenarios," and solved only *imprecisely*, to within a tolerance chosen *adaptively*, by balancing the estimated statistical error against solution error. The solutions from prior iterations serve as warm starts to aid efficient solution of the (piecewise linear convex optimization) sample-path problems generated on subsequent iterations. We characterize the almost sure and ℓ_1-norm convergence behavior. We also characterize the corresponding convergence rate; the latter rate is Monte Carlo canonical and analogous to the $\mathcal{O}(\epsilon^{-2})$ optimal complexity for non-smooth convex optimization. We report extensive large-scale numerical tests that clearly indicate favorable performance.

Green Simulation Optimization Using Likelihood Ratio Estimators

D. Eckman, Cornell University, United States, dje88@cornell.edu
B. Feng, University of Waterloo, Canada, ben.feng@uwaterloo.ca

Green simulation is the reuse of past simulation outputs to enhance the efficiency of current and future simulation experiments. One natural application of green simulation is in the context of simulation optimization, wherein outputs from past iterations in a search can be reused in subsequent iterations. In this article, we draw attention to challenges that arise when green simulation likelihood ratio estimators are naively employed in simulation optimization. In particular, we show that for searches that identify new designs based on past outputs, outputs in different iterations are conditionally dependent, violating one of the assumptions for the validity of the likelihood ratio estimator. As a result, green simulation likelihood ratio estimators of the objective and gradient can become biased. We demonstrate how this conditional dependence and bias can adversely affect the behavior of gradient-based optimization algorithms.

Parallel Adaptive Survivor Selection for VERY Large-Scale Simulation Optimization

B. Nelson, Northwestern University, United States, nelsonb@northwestern.edu
L. Pei, Northwestern University, United States, lindapei2016@gmail.com
S. Hunter, Purdue University, United States, susanhunter@purdue.edu

The class of methods known collectively as “ranking & selection” has been a theoretical and practical success story for finding the best among a finite number k of simulated systems (“feasible solutions”) with a statistical guarantee of correctness. However, the tricks that made these methods efficient in terms of the number of simulated realizations required can also render them (sometimes horribly) inefficient and invalid when implemented in a parallel, high-performance computing environment, which is of course exactly what one wants to do. Parallel adaptive survivor selection (PASS) jettisons the standard framework of ranking & selection to produce a new class of algorithms that *benefit* statistically from large k without any increase in coordination among parallel processors or added conservatism. Basic theory of PASS and the results of large-scale experiments will be presented.

Session 1.2 in Room P3 – Stochastic Networks and Scaling Limits
Chair: Amarjit Budhiraja

The divergence between Markovian and non-Markovian queueing models

R. Atar, Technion, Israel, rami@technion.ac.il
A. Budhiraja, University of North Carolina, USA, amarjit@unc.edu
P. Dupuis, Brown University, USA, paul_dupuis@brown.edu
R. Wu, University of Michigan, USA, ruoyu@umich.edu

Large deviation bounds are often very sensitive to the underlying distributions. An approach based on Renyi divergence allows one to address this issue by introducing bounds that are uniform within families of models specified in terms of this divergence. The talk will focus on applications of this approach to queueing models. For these models, key elements are the ability to (a) reduce the problem to that of Renyi divergence estimates merely on renewal processes, and (b) base the calculation on our better understanding of Markovian models. Examples to be described include the generalized Jackson network and the many server G/G/n+G.

Convergence rates to stationarity for reflecting Brownian motions

S. Banerjee, University of North Carolina, Chapel Hill, sayan@email.unc.edu
A. Budhiraja, University of North Carolina, Chapel Hill, amarjit@unc.edu

In this talk, we will explore methods to obtain rates of convergence to stationarity in L^1-Wasserstein distance for obliquely reflected d-dimensional reflected Brownian motion (RBM) in the nonnegative orthant that are explicit in the dimension and the system parameters. These processes emerge as heavy-traffic scaling limits of Jackson networks and convergence rates are essential in quantifying the efficiency of these continuous approximations. The results are then applied to a class of RBMs considered by Blanchet and Chen (2016) and to rank-based diffusions including the Atlas model. In both cases, we obtain explicit rates and bounds on relaxation times. In the first case we improve the relaxation time estimates of $O(d^4(\log d)^2)$ obtained in Blanchet and Chen (2016) to $O((\log d)^2)$. In the latter case, we give the first results on explicit parameter and dimension dependent rates under the Wasserstein distance. The proofs do not require an explicit form for the stationary measure or reversibility of the process with respect to this measure, and cover settings where these properties are not available. In the special case of the standard Atlas model, we obtain a bound on the relaxation time of $O(d^6(\log d)^2)$.

Asymptotic Behavior of a Critical Fluid Model for a Multiclass Processor Sharing Queue via Relative Entropy

J. A. Mulvany, University of Southern California, USA, justin.mulvany@marshall.usc.edu
A. L. Puha, California State University San Marcos, USA, apuha@csusm.edu
R. J. Williams, University of California San Diego, USA, williams@math.ucsd.edu

Queueing systems operating under the processor sharing discipline are relevant for studying time-sharing in computer and communication systems. Measure-valued processes, which track the residual service times of all jobs in the system, have been used to describe the dynamics of such systems. However, exact analysis of these infinite-dimensional stochastic processes is rarely possible. As a tool for approximate analysis of such systems, it has been proved that a fluid model arises as a functional law of large numbers limit of a multi-class processor sharing queue. This talk will focus on the asymptotic behavior of such a fluid model in the interesting regime of critical loading, where the average inflow of work to the system is equal to the capacity of the system to process that load.
Using an approach involving a certain relative entropy functional, we show that critical fluid model solutions converge to a set of invariant states as time goes to infinity, uniformly for all initial conditions lying in certain relatively compact sets. This generalizes an earlier single-class result of Puha and Williams to the more complex multiclass setting. In particular, several new challenges are overcome, including formulation of a suitable relative entropy functional and identifying a convenient form of the time derivative of the relative entropy applied to trajectories of critical fluid model solutions.

Reflected Diffusions and (Bio)Chemical Reaction Networks

D. F. Anderson, University of Wisconsin-Madison, USA, anderson@math.wisc.edu
D. J. Higham, University of Edinburgh, UK, d.j.higham@strath.ac.uk
S. C. Leite, Universidade Federal do ABC, Santo André, Brazil, saul.leite@ufabc.edu.br
R. J. Williams, University of California, San Diego, USA, williams@math.ucsd.edu

Continuous-time Markov chain models are often used to describe the stochastic dynamics of networks of reacting chemical species, especially in the growing field of systems biology. Discrete-event stochastic simulation of these models rapidly becomes computationally intensive. Consequently, more tractable diffusion approximations are commonly used in numerical computation, even for modest-sized networks. However, existing approximations (e.g., linear noise and Langevin), do not respect the constraint that chemical concentrations are never negative.
In this talk, we propose an approximation for such Markov chains, via reflected diffusion processes, that respects the fact that concentrations of chemical species are non-negative. This fixes a difficulty with Langevin approximations that they are frequently only valid until the boundary of the positive orthant is reached. Our approximation has the added advantage that it can be written down immediately from the chemical reactions. Some numerical examples illustrate the advantages of our approximation over direct simulation of the Markov chain or use of the linear noise approximation.

Session 1.3 in Room P4 – Job Scheduling and Load Balancing in Queueing Systems
Chair: Onno Boxma & Jan-Pieter Dorsman

Job Allocation in Large-Scale Service Systems with Affinity Relations

E. Cardinaels, Eindhoven University of Technology, The Netherlands, e.cardinaels@tue.nl
S. C. Borst, Eindhoven University of Technology, The Netherlands, s.c.borst@tue.nl
J. S. H. van Leeuwaarden, Eindhoven University of Technology, The Netherlands, j.s.h.v.leeuwaarden@tue.nl

Well-known load balancing algorithms, such as Join-the-Shortest-Queue (JSQ) or Join-the-Idle-Queue (JIQ), have been extensively analyzed in an overarching framework called the supermarket model. This model consists of a single dispatcher which directs the arriving jobs to one of the N identical parallel servers. A key feature of the supermarket model is the exchangeability of the servers in the sense that any job can be handled equally well by any server, which is often not the case in practice. We will focus on a scenario where jobs are not intrinsically different, but where particular servers might be better equipped to process certain jobs because of affinity or compatibility relations. These job-server affinity relations can model network topologies based on geographical proximity, or data locality in cloud scenarios. We utilize stochastic coupling techniques to obtain stability conditions and performance upper bounds for the introduced load balancing scheme. Moreover, a fluid limit analysis for symmetric model instances reveals a delicate interplay between the model parameters and the load balancing performance.

Replicate to the shortest queues

R. Atar, Technion, Israel, rami@ee.technion.ac.il
I. Keslassy, Technion, Israel, isaac@ee.technion.ac.il
G. Mendelson, Technion, Israel, galmen@campus.technion.ac.il

We introduce a load balancing policy that interpolates between two well known policies, namely *join the shortest queue* (JSQ) and join the least workload (JLW), and study it in heavy traffic. This policy, which we call *replicate to the shortest queues* (RSQ(d)), routes jobs from a stream of arrivals into buffers attached to N servers by replicating each arrival into $1 \leq d \leq N$ tasks and sending the replicas to the d shortest queues. When the first of the tasks reaches a server, its $d-1$ replicas are canceled. Clearly RSQ(1) is equivalent to JSQ, and it has been shown that RSQ(N) is equivalent to JLW; intermediate values of d provide tradeoff between good performance measures of JSQ and those of JLW. In heavy traffic, a key property underlying asymptotic analysis of load balancing policies is *state space collapse* (SSC). Unlike policies such as JSQ, where SSC is well-understood, the treatment of SSC under RSQ(d) requires to address the massive cancellations that highly complicate the queue length dynamics. Our first main result is that SSC holds under RSQ(d) for possibly heterogeneous servers. Based on this result we obtain diffusion limits for the queue lengths in the form of one-dimensional reflected Brownian motion, asymptotic characterization of the short-time-averaged delay process, and a version of Reiman's snapshot principle. We illustrate using simulations that, as d increases, the server workloads become more balanced, and the delay distribution's tail becomes lighter. We also discuss the implementation complexity of the policy as compared to that of the *redundancy routing* policy, to which it is closely related.

Achievable stability in redundancy scheduling

Y. Raaijmakers, Eindhoven University of Technology, The Netherlands, y.raaijmakers@tue.nl
S. C. Borst, Eindhoven University of Technology, The Netherlands, s.c.borst@tue.nl
O. J. Boxma, Eindhoven University of Technology, The Netherlands, o.j.boxma@tue.nl

We consider a system with N parallel servers, where a dispatcher immediately assigns an arriving job to one or more servers. We investigate the stability region for the 'cancel-on-completion' version of redundancy-d in a quite general workload model, i.e., we allow for generally distributed job sizes and the speeds for a given job type are allowed to be inter-dependent and non-identically distributed. Under the assumption that job types are beforehand known and d identical replicas are assigned to the servers according to a static probabilistic job allocation, we show that no replication $(d = 1)$ gives a strictly larger stability region than replication $(d > 1)$. Furthermore, we show that for independent and identically distributed (i.i.d.) replicas the same statement is true under the additional assumption of a New-Better-than-Used (NBU) job size distribution. Both for identical and i.i.d. replicas observing the job types is essential for these statements, since there are scenarios in which full replication $(d = N)$ outperforms no replication $(d = 1)$ when job types can not be observed.
Numerical results are presented to quantify the performance loss when job types are beforehand unknown, but where we could learn the job type. Information about the job type can (only) be obtained while the job is in service. During this service we allow for a job to be transferred (or replicated). More specifically, if the job is in service and has still not been completed, the belief that the server speed of this job is slow increases, and we could try serving this job on different (possibly faster) servers.

Towards optimality in parallel job scheduling

B. Berg, Carnegie Mellon University, Pittsburgh, PA, United States of America, bsberg@cs.cmu.edu
J. L. Dorsman, University of Amsterdam, The Netherlands, j.l.dorsman@uva.nl
M. Harchol-Balter, Carnegie Mellon University, Pittsburgh, PA, United States of America, harchol@cs.cmu.edu

Running jobs in parallel is an excellent way to reduce their response time. To this end, modern multi-core chips afford the opportunity to reduce job response time by parallelizing a job across several cores. This leads to the question of how many cores to assign to each job, also called the level of parallelization. In making this choice, one is confronted with the following trade-off. Parallelizing an individual job across multiple cores reduces the response time of that individual job. In practice, however, the speedup a job experiences is sublinear and concave in its level of parallelization, leading to an inefficient use of resources and additional system load. Hence, while a higher level of parallelization may decrease an individual job's response time, it may have a deleterious effect on overall response time.
In this presentation, we will study this trade-off. In particular, when all jobs share the same size distribution and the same concave speedup curve, we will argue that using a fixed level of parallelization is near-optimal. We will also consider the case where jobs may follow different speedup curves, in which case finding a good scheduling policy is even more challenging.

Session 1.4 in Room P5 – Stochastic Systems
Chair: Mark Squillante & David Yao

Why is Kemeny's constant a constant?

D. Bini, University of Pisa, Italy, bini@dm.unipi.it
J. Hunter, Auckland University of Technology, New Zealand, jeffrey.hunter@aut.ac.nz
G. Latouche, Université Libre de Bruxelles, Belgium, latouche@ulb.ac.be
B. Meini, University of Pisa, Italy, beatrice.meini@unipi.it
P. Taylor, University of Melbourne, Australia, taylorpg@unimelb.edu.au

In their 1960 book on finite Markov chains, Kemeny and Snell established that a certain sum is invariant. The value of this sum has become known as *Kemeny's constant*. Various proofs for the invariance have been given over time, some more technical than others. We shall first give a simple algebraic proof and then follow it up with a probabilistic proof that gives physical insight into what is going on. The result extends without a hitch to continuous-time Markov chains on a finite state space.
For Markov chains with denumerably infinite state space, the constant may be infinite and even if it is finite, there is no guarantee that the physical argument will hold. We shall show that the physical interpretation does go through for the special case of a birth-and-death process with a finite value of Kemeny's constant.

Unique Ergodicity of Diffusion Limits of Many-Server Queues with Reneging

P. Agarwal, Brown University, United States of America, pooja_agarwal@brown.edu
K. Ramanan, Brown University, Finland, United States of America, kavita_ramanan@brown.edu

We consider the so-called $GI/GI/N+GI$ queue, which is a queueing system with N servers in which impatient customers with i.i.d. service times, and i.i.d. patience times enter service in the order of arrival and abandon the queue if the time before possible entry into service exceeds the patience time. Such systems arise as models in diverse applications, ranging from operations research to biochemical reaction networks. Since an exact analysis is infeasible, we consider an asymptotic analysis as the number of servers goes to infinity, which is the regime of interest in many applications. Under suitable assumptions, we establish a functional central limit theorem, and characterize the limit as the unique solution to a non-standard stochastic partial differential equation, subject to a non-linear boundary condition. We also show that the limit process describing the total number in system is an Itô diffusion. A measure-valued representation of the state process plays an important role in the analysis. We also show that the infinite-dimensional limit process has a unique invariant measure.

Exponential Holding Cost in a Single-Server Queue: Interchange of Limits and Asymptotic Optimality of Reflection Control

J. Yang, School of Science, Beijing University of Posts and Telecommunications, Beijing, China, yangjk@amss.ac.cn
D. D. Yao, Dept of Industrial Engineering and Operations Research Columbia University, New York, USA, yao@columbia.edu
H. Q. Ye, Faculty of Business, Hong Kong Polytechnic University, Hong Kong, China, hq.ye@polyu.edu.hk

Consider base-stock control in a production-inventory system modeled as a single-server queue, GI/GI/1. The work-in-process inventory on the input side often consists of raw materials, and as such requires financing one way or another; i.e., the inventory should be priced similar to a financial asset. This motivates us to consider an exponential holding cost function, or equivalently, the exponential moments of the GI/GI/1 model. We study the asymptotic optimality of the base-stock control (or "reflection control") in this setting, and the associated interchange of limits (with respect to time and to the scaling constant).

Fatou's lemma in its classic form for varying measures with applications to MDPs

E. A. Feinberg, Stony Brook University, USA, eugene.feinberg@stonybrook.edu
P. O. Kasyanov, National Technical University of Ukraine, Ukraine, kasyanov@i.ui
Y. Liang, University of Toronto, Canada, yan.liang@rotman.utoronto.ca

The classic Fatou lemma states that the lower limit of a sequence of integrals of functions is greater or equal than the integral of the lower limit. It is known that Fatou's lemma for a sequence of weakly converging measures states a weaker inequality because the integral of the lower limit is replaced with integral of the lower limit in two parameters, where the second parameter is the argument of the functions. This paper provides sufficient conditions when Fatou's lemma holds in its classic form for a sequence of weakly converging measures. The functions can take both positive and negative values. The paper also provides similar results for sequences of setwise converging measures. It also provides Lebesgue's and monotone convergence theorem for sequences of weakly and setwise converging measures. The obtained results are used to prove broad sufficient conditions for the validity of optimality equations for average-costs Markov decision processes.

Session 1.5 in Room M1 – Statistical Methods for Learning and Decision-Making
Chair: Nan Ye

Comparison of machine learning algorithms in classification of grazing behaviour in sheep

S. Hu, Queensland University of Technology, Brisbane, shuwen.hu@hdr.qut.edu.au
Y. Li, CSIRO, Brisbane, yutao.li@csiro.au
A. Ingham, CSIRO, Brisbane, aaron.ingham@csiro.au
G. B. Hurley, CSIRO, Brisbane, greg.bishop-hurley@csiro.au
E. G. Gonzalez, French National Institute for Agricultural Research, France, eliel.gonzalez-garcia@inra.fr
Y.-G. Wang, Queensland University of Technology, Brisbane, you-gan.wang@qut.edu.au

Grazing, or feed intake, is an important daily activity for ruminant livestock species as this behaviour controls the primary input variable into animal production systems. Inertial motion sensors located on the animal have been used to study the grazing behaviour of ruminant livestock where patterns in the sensor signal can be correlated to the observed behaviours. However, to date most research on automatic livestock behaviour has been focused on cattle. In this talk, we describe the application of three machine learning algorithms (Random Forests (RF), Support Vector Machine (SVM) and Linear Discriminate Analysis (LDA)) to classify grazing behaviour patterns in inertial sensor data collected from collars worn on the neck of sheep. A total of six Romney ewe lambs that carried sensors generating time series accelerometer data was used in the study. Initially the individual sensor data of the animals and the group behaviour information from 15 lambs were consolidated for the evaluation of the utility of three machine learning methods. Six behaviours were recorded, including grazing, ruminating, standing, walking, resting and other (a collection of all other behaviours). A range of statistical features (e.g. mean, variance, energy, entropy) were examined for the identification of the suitable classification algorithm that can be used to accurately estimate the total grazing time of individual animals. The results were then combined with true measures of biomass disappearance to model grass intake with the aim to identify the most feed efficient animals that can manage available biomass within the landscape. The outcomes from the comparisons of different methods and feature selections will be presented.

Efficient Newton-type Methods for Large-scale Invex Optimization

Y. Liu, University of Queensland, Australia, yang.liu2@uq.edu.au
F. Roosta, University of Queensland, Australia, and International Computer Science Institute, Berkeley, USA, fred.roosta@uq.edu.au

We consider the application of second-order optimization methods, namely variants of the classical Newton's method, for invex problems. Motivated by our prior work on Newton-MR, we here introduce an efficient inexact variant of Newton-MR for large-scale invex problems, called stochastic Newton-MR, in which the Hessian information is stochastically approximated. Our work relies on some ideas from matrix projections and perturbation theory. In particular, as a key step in the convergence proof of stochastic Newton-MR, given a bound on the spectral norm of the perturbation in the Hessian, we derive a similar bound related to Hessian pseudo-inverses. From this, we are able to derive linear convergence rates for stochastic Newton-MR. Experiments show a great promise for our algorithm.

Chaotic time series regression modeling using phase space reconstruction and deep neural network

Y. G. Wang, Queensland University of Technology, Australia, you-gan.wang@qut.edu.au
J. Wu, Queensland University of Technology, Australia, jinran.wu@hdr.qut.edu.au

The conventional time series modeling focuses on the simple linear time series; however, most of time series are extremely complex and non-linear dynamic systems in real world. To forecast these series with high precision, we propose a novel hybrid forecasting model, which combines phase space reconstruction and deep neural network learning. With phase space reconstruction we calculate the embedding dimension and optimal delay time for the chaotic and non-linear time series , hence determining more appropriate features for deep learning model. The deep neural network made robust by cross validation then uses these intelligently-selected features to provide much improved forecasting. We demonstrate the superiority of our new model by applying it to a case study of Australian electricity scheduling, using half-hour data collected from New South Wales, Queensland and Victoria, Australia (January 1, 2016 - March 20, 2019). Our model outperforms all of our benchmark approaches in time series forecasting.

Maximum Entropy Approaches for Inverse Reinforcement Learning

A. J. Snoswell, The University of Queensland, Australia, a.snoswell@uq.edu.au
S. P. N. Singh, The University of Queensland, Australia, spns@uq.edu.au
N. Ye, The University of Queensland, Australia, nan.ye@uq.edu.au

We make decisions to maximize our perceived reward, but handcrafting a reward function for an autonomous agent is challenging. Inverse Reinforcement Learning (IRL), which is concerned with learning a reward function from expert demonstrations, has recently attracted significant interest, with the Maximum Entropy (MaxEnt) approach being a popular method.

In this talk, we will explore and contrast a variety of MaxEnt IRL approaches. We show that in the presence of stochastic dynamics, a minimum KL-divergence condition provides a rigorous derivation of the MaxEnt model, improving over a prior heuristic derivation. Furthermore, we explore extensions of the MaxEnt IRL method to the case of unknown stochastic transition dynamics, including a generative model for trajectories, a discriminative model for action sequences, and a simple logistic regression model.

We will present evaluation results for simulated and real-world problems, including the UCI Taxi Service Trajectory dataset, which considers the problem of long-distance forecasting of driver behaviour.

Session 1.6 in Room M2 – Rare Events and Distributionally Robust Optimization
Chair: Chang-Han Rhee

Confidence Regions for Optimal Transport Distributionally Robust Optimization

J. Blanchet, Stanford University, USA, jose.blanchet@stanford.edu
K. Murthy, SUTD, Singapore, karthyek_murthy@sutd.edu.sg
N. Si, Stanford University, USA, niansi@stanford.edu

We establish the asymptotic normality of estimators obtained out of solving data-driven distributionally robust optimization with an uncertainty region specified in terms of the Wasserstein distance. Moreover, we characterize an asymptotically optimal confidence region in terms of convergence of sets (in the so-called Fell topology).

Safety evaluation of black-box prediction models via rare-event simulation

Z. Huang, University of Michigan, United States, zhyhuang@umich.edu
H. Lam, Columbia University, United States, henry.lam@columbia.edu
D. Zhao, Carnegie Mellon University, United States, dingzhao@cmu.edu

We study the design of good importance samplers to simulate the probability that a black-box predictor, built for instance from machine learning tools such as neural networks and random forests, outputs an exceedingly high prediction value. This problem is motivated from robustness and safety estimations that arise in recent applications of autonomous vehicles and other physical systems. Our approach utilizes mixed integer programming and a “cutting plane” approach to sequentially locate dominant points that guide the tilting of distributions in the importance samplers.

Sample path large deviation Levy processes and random walks with regularly varying increments in multiple dimensions

Z. Su, Northwestern University, USA, zhesu2022@u.northwestern.edu
C.-H. Rhee, Northwestern University, USA, chang-han.rhee@northwestern.edu

In many applications, it has been observed that there is a structural difference in the way system-wide rare events arise when the underlying uncertainties are heavy-tailed (compared to the better-understood light-tailed counterparts). Roughly speaking, in light-tailed settings, the system-wide rare events arise because everything goes wrong a little bit (conspiracy principle), whereas in heavy-tailed settings, the system-wide rare events arise because only a few things go wrong, but they go catastrophically wrong (catastrophe principle). Due to such a fundamental difference, as well as the ubiquitous presence of the heavy-tailed distributions in modern engineering systems, a comprehensive theory of large deviations for heavy-tailed rare events has long been called for.

In the past couple decades, there have been a few breakthroughs in understanding the catastrophe principle. In particular, Hult, Lindskog, Mikosch, and Samorodnitsky (2005) was the first systematic confirmation of the catastrophe principle for random walks with multi-dimensional regularly varying increments in case the rare-events of interest can be driven by a single unusually large increment (a.k.a. principle of a single big jump). More recently, for 1-dimensional processes, Rhee, Blanchet, and Zwart further characterized the catastrophe principle in the form of sample path large deviations for a general class of rare-events far beyond the principle of a single big jump—i.e., the rare-events that can be caused by an arbitrary number of large increments.

In this talk, we establish the sample path large deviations for random walks and Lévy processes with multi-dimensional regularly varying increment distributions and for rare-events that are driven by an arbitrary number of large jumps. This completely characterizes the catastrophe principle and generalizes known results.

Catastrophe principle for correlated stochastic processes

C.-H. Rhee, Northwestern University, USA, chang-han.rhee@northwestern.edu
M. Bazhba, CWI, Netherlands, M.Bazhba@cwi.nl
B. Chen, CWI, Netherlands, B.Chen@cwi.nl
J. Blanchet, Stanford University, USA, jose.blanchet@stanford.edu
B. Zwart, CWI, Netherlands, Bert.Zwart@cwi.nl

Heavy-tails are observed in many man-made systems. Examples include financial losses, income distributions, degree distributions in social networks, CPU requirements and file sizes in computer systems. Heavy-tailed distributions have many strange and beautiful properties, but the one that is particularly relevant to the study of stochastic systems is the *catastrophe principle*: when underlying uncertainties are heavy-tailed, system-wide rare events arise due to the catastrophic failures of a small number of components. In stark contrast, light-tailed systems typically experience a system-wide rare event because all the components exhibit subtly different behaviors as if the entire system has conspired to provoke the rare-event. Recent advances in heavy-tailed large deviations theory have successfully characterized the catastrophe principle in rigorous mathematical terms for random walks and Lévy processes where the increment distributions are heavy-tailed. Together with the classical large deviations theory for light-tailed random walks and Lévy processes, the recent heavy-tailed large deviations results establishes a dichotomy in case the stochastic processes have independent increments. Under the presence of correlation, however, such dichotomy gets blurred. In particular, heavy-tailed behavior can arise—for example, due to a multiplcative mechanism—even when each individual uncertainties are light-tailed. In such contexts, more advanced tools and subtle treatments are required to understand the system-wide rare events. In this talk, we present our new theoretical developments for dealing with such cases and illustrate the treatment of correlated processes with Markov random walks driven by two stochastic processes that are important in operations research—Lindley's recursion and auto-regressive processes.

Session 1.7 in Rooms M5 and M6 – Control and Simulation of Stochastic Systems Chair: Jiheng Zhang

Perfect Sampling for Queues with Autoregressive Arrivals

X. Chen, Chinese University of Hong Kong, Shenzhen, China, chenxinyun@cuhk.edu.cn

Motivated by the applications in social media, healthcare and financial systems, queueing models with autoregressive arrivals, such as the Hawkes or shot-noise-driven arrivals, have been studied in several recent works. So far, there is no analytic approach to compute the stationary distribution of single- or multiple - server queues. In this talk, we propose a perfect sampling algorithm that generates i.i.d. samples exactly following the stationary distribution of single-server queues with autoregressive arrivals.

Dynamic Routing in a Many-Server System

P. Cao, University of Science and Technology of China, pcao@ustc.edu.cn
Z. Zhong, The Chinese University of Hong Kong, zhongzhiheng@baf.cuhk.edu.hk
J. Huang, The Chinese University of Hong Kong, junfeih@cuhk.edu.hk

We consider a service system with multiple stations and a single class of customers. Each station has many servers and its dedicated queue. We propose a minimum-expected-delay faster-server-first with ξ-choice (named MED-FSF(ξ)) routing policy, under which a subset of ξ (a discrete random variable) stations will be selected. The system manager then routes a new customer to one of the selected stations following the MED-ESF policy. We prove that our proposed routing policy is asymptotically equivalent to the MED-FSF policy. Simulations are used to to validate the results.

Dynamic Scheduling of Multiclass Many-server Queues with Abandonment: the Generalized $c\mu/h$ Rule

Z. Long, Hong Kong University of Science and Technology, Hong Kong S.A.R., China, zlong@ust.hk
N. Shimkin, Technion, Haifa, Israel, shimkin@ee.technion.ac.il
H. Zhang, Hong Kong University of Science and Technology, Hong Kong S.A.R., China, hzhangaq@connect.ust.hk
J. Zhang, Hong Kong University of Science and Technology, Hong Kong S.A.R., China, jiheng@ust.hk

We consider the problem of server scheduling in a multiclass many-server queueing system with abandonment. For the purpose of minimizing the long-run average queue length costs and abandon penalties, we propose three scheduling policies to cope with any general cost functions and general patience time distributions. First, we introduce the target-allocation policy, which assigns higher priority to customer classes with larger deviation from the desired allocation of the service capacity, and prove its optimality for any general queue length cost functions and patience time distributions. The $Gc\mu/h$ rule, which extends the well-known $Gc\mu$ rule by taking abandonment into account, is shown to be optimal for the case of convex queue length costs and nonincreasing hazard rates of patience. For the case of concave queue length costs but nondecreasing hazard rates of patience, it is optimal to apply a fixed priority policy, and a knapsack-like problem is developed to determine the optimal priority order efficiently. As a motivating example of the operations of emergency departments, a hybrid of the $Gc\mu/h$ rule and the fixed priority policy is suggested to reduce crowding and queue abandonment. Numerical experiments show that this hybrid policy performs satisfactorily.

An Asymptotically Optimal Index Policy for the Finite-Horizon Restless Bandit

W. Hu, Google, USA, smithweici@gmail.com
P. Frazier, Cornell University, USA, pf98@cornell.edu
X. Zhang, Cornell University, USA, xz556@cornell.edu

We consider the restless bandit, a generalization of the multi-armed bandit in which arms may change state when they are not pulled. In the stochastic infinite-horizon setting, Whittle proposed the Whittle index policy and conjectured the per-arm optimality gap vanishes as the number of arms grows to infinity while holding fixed the fraction of arms that can be pulled per period. Weber and Weiss showed this conjecture is true when arms have only 3 states, or when the fluid limit has a globally stable equilibrium point. We propose a novel but related index policy for the finite-horizon setting and show it is asymptotically optimal in the same sense without restrictions on the fluid limit. This policy does not require the challenging-to-verify indexability condition required by the Whittle index policy. It also generalizes naturally to weakly-coupled dynamic programs where it retains asymptotic optimality.

Session 1.8 in Rooms M7 and M8 – Stochastic Operations Research for Power Grids Applications
Chair: Alessandro Zocca

Learning Graph Parameters from Linear Measurements: Fundamental Trade-offs and Application to Electric Grids

T. Li, California Institute of Technology, United States, tongxin@caltech.edu
L. Werner, California Institute of Technology, United States, lwerner@caltech.edu
S. H. Low, California Institute of Technology, United States, slow@caltech.edu

We consider a specific graph learning task: reconstructing a symmetric matrix that represents an underlying graph using linear measurements. We study fundamental trade-offs between the number of measurements (sample complexity), the complexity of the graph class, and the probability of error by first deriving a necessary condition (fundamental limit) on the number of measurements. Then, by considering a two-stage recovery scheme, we give a sufficient condition for recovery. Furthermore, assuming the measurements are Gaussian IID, we prove upper and lower bounds on the (worst-case) sample complexity. In the special cases of the uniform distribution on trees with n nodes and the Erdős-Rényi (n, p) class, the fundamental trade-offs are tight up to multiplicative factors. Applying the Kirchhoff's matrix tree theorem, our results are extended to the scenario when part of the topology information is known a priori. In addition, we design and implement a polynomial-time (in n) algorithm based on the two-stage recovery scheme. Simulations for several canonical graph classes and IEEE power system test cases demonstrate the effectiveness of the proposed algorithm for accurate topology and parameter recovery.

Optimal Switching Control of Continuous-Time Stochastic Systems: Theory and Applications

R. Martyr, Queen Mary University of London, United Kingdom, r.martyr@qmul.ac.uk

The problem of optimal switching, or alternating control, features a decision maker, the controller, that controls a dynamical system over time by sequentially choosing an operating mode, or regime, from a finite set of such modes. The controller has an objective function that measures the performance of an alternating control, which includes a running reward, a cost for each time the operating mode is switched, and a terminal reward. The controller's objective is to maximise this objective function over all feasible alternating controls. Optimal switching problems often arise in the analysis of flexible industrial projects (Real Options Analysis), with examples related to investment in electricity generation and valuation of energy storage. A typical example is the economic valuation of a power plant which manages its fuel mix for electricity production. This talk overviews the theory and aforementioned applications of optimal switching problems, focusing on continuous-time stochastic systems.

No need to rush: dealing with deadlines in EV charging

A. Ferragut, Universidad ORT Uruguay, Uruguay, ferragut@ort.edu.uy
M. Zeballos, Universidad ORT Uruguay, Uruguay, mzeballos@uni.ort.edu.uy
F. Paganini, Universidad ORT Uruguay, Uruguay, paganini@ort.edu.uy

In this work, we consider a queueing system where tasks arrive as a Poisson process and have generally distributed job sizes. Each task has a (generally distributed) deadline, correlated with the job length. The system may serve each task k at any rate $0 \leqslant r_k \leqslant 1$, and has constrained capacity, i.e. $\sum_k r_k \leqslant C$ where the sum is over simultaneous jobs. These leave the system either when service is complete or upon expiration of its deadline. Inspired by the electric vehicle charging facilities, we consider however that *partial service* performed on each job is useful, since the energy will be available to the EV.
In this scenario, we are interested on how the service is distributed across users. Building upon previous formulations for the processor-sharing policy (Gromoll et al. (2006), Aveklouris et al. (2017), we propose a fluid limit analysis for such systems that includes many common policies such as earliest-deadline-first, least-laxity first or shortest-remaining-processing time. We show that in the limit the reneged work is invariant across efficient policies, but its distribution is highly dependent on policy. This leads to potential unfairness across jobs. We then derive a natural policy, dubbed least-laxity-ratio, that preserves proportional fairness in overload. Finally, we extend our modeling formulation to include age-based policies such as FIFO, LIFO or Least Attained Service. Such age based policies are interesting to deal with the case where the system does not know the deadline. We show that in overload, LIFO, Least-attained-service and Earliest Deadline First have exactly the same performance: a striking result given that the latter makes explicit use of the deadlines. Similar analogies can be built between other deadline-aware, age-aware policies and size-aware policies.
This work was partially supported by ANII Uruguay through grants FCE_1_2017_1_136748 and FSE_1_2016_1_131605.

Less is More: Failure Localization in Power Systems

A. Zocca, California Institute of Technology, azocca@caltech.edu
L. Guo, California Institute of Technology, lguo@caltech.edu
C. Liang, California Institute of Technology, cliang2@caltech.edu
S. H. Low, California Institute of Technology, slow@caltech.edu
A. Wierman, California Institute of Technology, adamw@caltech.edu

Cascading failures in power grids exhibit non-local propagation patterns which make the analysis and mitigation of failures difficult. A new approach to address such challenges has emerged in recent years using the notion of network tree-partition and has yielded many interesting, sometimes counter-intuitive, insights for the planning and management of power grids. For instance, we prove that, by properly reducing system redundancy, one can improve system robustness against component failures. In this talk I will present this approach in more detail and show how, given a power grid topology and flow pattern on it, one can identify the best partition (using ideas from the community detection literature) and optimally choose the lines to switch off to improve the system reliability.

Wednesday 13:15 – 14:45

Session 2.1 in Rooms P1 and P2 – Model Uncertainty and Simulation Optimization
Chair: Enlu Zhou

Sequential risk set inference for simulation optimization under input uncertainty

E. Song, Pennsylvania State University, United States, eus358@psu.edu

Simulation optimization (SO) is subject to model risk when the input distributions of the simulation model are estimated based on finite real-world data. We discuss a different way of looking at a SO problem under model risk by introducing the practical tolerance parameter, $\delta > 0$, specified by the user. We define an α-level risk set as a set of solutions that perform better than the conditional best given the best estimate of the input distributions by a margin greater than δ with probability greater than α, where the probability is taken with respect to the distribution over the candidate input distributions. Taking a Bayesian approach, we define a posterior distribution over the candidate input distributions conditional on the real-world observations. We model the conditional means of solutions given input distributions with a Gaussian process (GP) prior and update the GP based on the cumulative simulation results. The α-level risk set is estimated from the posterior distribution of GP as well as the posterior on the input distributions. For efficient estimation of the risk set, we propose two sequential sampling criterion that selects a solution-input distribution pair to simulate at each iteration to minimize the estimation error of the risk set using the predictive distribution of the GP.

Computationally Efficient Quantification of Simulation Input Uncertainty

H. Lam, Columbia University, New York, US, henry.lam@columbia.edu
H. Qian, Columbia University, New York, US, h.qian@columbia.edu

Quantifying the input uncertainty of stochastic simulation, based on variance estimation or confidence interval construction, often entails substantial simulation costs due to the nested sampling requirements of the quantification schemes. We present several related approaches to reduce these costs, utilizing optimization, subsampling, and random perturbation respectively. We explain the statistical mechanisms of these approaches and why they need significantly less simulation efforts than some previous methods. We also compare them in terms of the ease of implementation and empirical performances.

On the convergence rates of simulation with covariates

C. Li, National University of Singapore, Singapore, stalic@nus.edu.sg
S. Gao, City University of Hong Kong, Hong Kong, siyangao@cityu.edu.hk

We consider performing simulation experiments in the presence of covariates. Here, covariates refer to some input information other than system designs to the simulation model which also affects the system performance. Simulation with covariates is a recently proposed framework, which advocates starting the simulation offline before the covariate values are revealed and using metamodeling techniques to predict the desired results. This framework is able to significantly reduce the computing time compared the traditional online simulation. In this paper, we follow this framework, adopt the stochastic kriging (SK) prediction model and study how fast the prediction errors diminish with the size of samples collected in the covariate space. This is a fundamental problem in simulation with covariates and helps quantify the relationship between the offline simulation efforts and the online prediction accuracy. Particularly, we adopt the maximal integrated mean squared error (IMSE) and integrated probability of false selection (IPFS) as measures for assessing the pre- diction errors for the purposes of system performance evaluation and optimization respectively. Then, we establish convergence rates for the two measures under different covariance kernels and conditions. Last, these convergence behaviors are illustrated numerically using test examples.

Online Quantification of Input Uncertainty for Parametric Models

E. Zhou, Georgia Institute of Technology, enlu.zhou@isye.gatech.edu
T. Liu, Georgia Institute of Technology, tianyiliu@gatech.edu

It has become increasingly important to assimilate "online data" that arrive sequentially in time for real-time decision. Input uncertainty quantification in stochastic simulation has been developed extensively for batch data that are available all at once, but little has been studied for online data. We propose a computationally efficient method to incorporate online data in realtime for input uncertainty quantification of parametric models. We show finite-sample bounds and asymptotic convergence for the proposed method, and demonstrate its performance on a simple numerical example.

Session 2.2 in Room P3 – Queueing Networks and Markov Chains
Chair: Jim Dai

State Space Collapse for Multi-class Queueing Networks with SBP Service Policies

C. Cao, Cornell University, USA, cc2488@cornell.edu
J. G. Dai, Cornell University, USA, jd694@cornell.edu
X. Y. Zhang, Cornell University, USA, xz556@cornell.edu

We introduce the notion of strong state space collapse(SSSC) for the high priority classes in multiclass queueing networks(MCN) and fluid state space collapse(FSSC) for the corresponding fluid models. SSSC holds for a sequence of MCNs in heavy traffic if the pth moments of the stationary queue length are uniformly bounded. FSSC holds for the corresponding sequence of fluid models if they uniformly drain to zero, including the critical fluid model.
We prove that SSSC holds for a sequence of multiclass queueing networks in heavy traffic under static buffer priority(SBP) service discipline if FSSC holds. We will also provide examples of Lyapunov functions to illustrate how one can check the FSSC condition.

Accelerating simulated annealing via replica exchange

M. Choi, The Chinese University of Hong Kong, Shenzhen, China, michaelchoi@cuhk.edu.cn

In classical simulated annealing, the algorithm constructs a non-homogeneous Markov chain that converges to the set of global minima of the target function as the temperature cools down gradually. The "optimal" cooling schedule is of the form $\frac{c}{\ln(t+1)}$, where c is known as the hill-climbing constant as proved in Hajek (1988). In this talk, we demonstrate how we can speed up simulated annealing using the idea of replica exchange. The core idea is to run two simulated annealing algorithms in parallel with the possibility of swapping their configurations.

Steady-state heavy traffic limits for multiclass queueing networks with SBP service policies

A. Braverman, Northwestern University, U.S.A., anton.braverman@gmail.com,
J. G. Dai, Cornell University, U.S.A., jim.dai@gmail.com,
M. Miyazawa, Tokyo University of Science, Japan, miyazawa@rs.tus.ac.jp

We are interested in the stationary distribution of a multiclass queueing network with static buffer priority (SBP) service discipline, and consider its diffusion approximation under heavy traffic. Unlike the standard approach using limit interchange, we directly derive the diffusion approximation through basic adjoint relationship (BAR) of the stationary distribution. For this, we assume state space collapse (SSC) in distribution in addition to the stability of prelimit processes and the usual heavy traffic condition. In this talk, we focus on how SSC assumption arises in our BAR approach and how it is used in deriving the diffusion limit of the stationary distribution. Sufficient conditions for this SSC assumption are provided by Chang Cao and Jim Dai in the same session.

Session 2.3 in Room P4 – Strategic Behavior in Queues
Chair: Moshe Haviv

Coffee Shop Operations with Mobile Ordering

K. Kang, University of Minnesota, US, kangx747@umn.edu
S. Doroudi, University of Minnesota, US, sdoroudi@umn.edu
M. Delasay, Stony Brook University, US, mohammad.delasay@stonybrook.edu

Mobile ordering is becoming popular in the world of coffee shops, fast food, and sit-down restaurants. Our work examines the operations of a coffee shop where some customers can use a mobile app to skip the ordering and payment queue. Using queueing theory, we explore the impact of mobile customers on walk-in customers and vice versa across a variety of service policies. Our work seeks to understand the tradeoffs between helping different classes of customers by identifying the space of managerial policies that allow service capacity and priority to be distributed across these two customer classes. We also study the strategic behavior of walk-in customers under a partially observable tandem queueing system and numerically investigate the relationship between customer loss and customer adoption of mobile ordering in equilibrium.

Regulating Service Length Demand in a Single Server Queue

B. Oz, Hebrew Univesity of Jerusalem, Israel, b.oz@huji.ac.il

We consider an unobservable M/?/1 queue where costumers arrive according to a Poisson process and decide strategically on their own service length in order the maximize their welfare. Customers are heterogeneous with respect to their service valuation, that is increasing and concave in the service length, and are homogeneous with respect to their waiting costs, that are linear in the time spent in the system. In a symmetric equilibrium, the joining strategy is based on a threshold of the individual service valuation while the service length is increasing in the service evaluation. When comparing the equilibrium behavior to the socially optimal strategy, too many are joining the queue, and those who join overuse the server. We discuss a number of regulation methods to incentivise socially optimal behavior.

Equilibrium Threshold Strategy for An $M/M/1$ Feedback Queue

M. Fackrell, The University of Melbourne, Australia, fackrell@unimelb.edu.au
P. Taylor, The University of Melbourne, Australia, taylorpg@unimelb.edu.au
J. Wang, The University of Melbourne, Australia, jiesenw@student.unimelb.edu.au

We investigate customers' strategic behavior in an $M/M/1$ feedback queue with linear waiting cost and fixed reward. At their arrival time, customers can observe the length of the queue and then decide to join or to balk. Customers are served in order of arrival. After being served, each customer either departs the system with probability q, or immediately joins the back of the queue otherwise. For each customer, the feedback induces a different best response, which is a function of not only the queue length, but also the other customers' strategies. The equilibrium strategy is of threshold type: customers choose to join the queue when the queue length is shorter than some threshold value, and balk otherwise.
In particular, for a customer who joins at a specific position in the queue, we get an analytical expression for his/her expected sojourn time if all the other customers always choose to join the queue. We also compute the conditional expected sojourn time based on other customers' threshold values by solving Poisson's equation for a discrete time nonhomogeneous quasi-birth-and-death-process. Finally, we show via *coupling* that the expected sojourn time is non-decreasing with the initial queue size and other customers' threshold values, based on which, we propose our Nash equilibrium threshold strategy.

On the (sub)optimality of the $c\mu$ rule when customers are strategic

M. Haviv, Hebrew University of Jerusalem, Israel, moshe.haviv@gmail.com
B. Oz, Hebrew University of Jerusalem, Israel, b.oz@huji.ac.il

When a queueing system manager faces an heterogeneous population of customers the question of priorities naturally arises. In case that heterogeneity is with respect to waiting cost per unit time and social welfare is of interest, it is well known that the $c\mu$ rule is optimal as long as everything else is fixed. However, optimality is no longer guaranteed if the arrival process is endogenous due to strategic customer behavior. This leads to a schizophrenia towards customers whose cost of waiting per unit time is high. On one hand, it is sometimes socially best that they will not join at all as the gain in the throughput is overwhelmed by the added waiting costs. On the other hand, once they join, it is socially optimal to give them priority. Using a standard queueing model, we show that if an efficient admission control policy can be used to deter too costly customers from joining, the $c\mu$ rule is optimal. Otherwise, applying the $c\mu$ rule leads to an unusual equilibrium joining behavior. In particular, it is not a threshold-based one. Specifically, the equilibrium joining probability as a function of the cost parameter alternates between intervals where joining comes with some (cost dependent) probability strictly between zero and one, and intervals of joining with probability one. This leads to the conclusion that it is not a priory clear whether the $c\mu$ rule is socially better than the first-come first-served discipline when admission cannot be centrally controlled.

Session 2.4 in Room P5 – Asymptotic Analysis in Stochastic Systems
Chair: Guodong Pang

Limit theorems for record indicators in threshold Nevzorov schemes

K. Borovkov, The University of Melbourne, Australia, borovkov@unimelb.edu.au
P. He, The University of Melbourne, Australia, patrick.he95@gmail.com

In Nevzorov's F^{α}-scheme, one deals with a sequence of independent random variables whose distribution functions are all powers of a common continuous distribution function. A key property of the F^{α}-scheme is that the record indicators for such a sequence are independent. This allows one to obtain several important limit theorems for the total number of records in the sequence up to time $n \to \infty$. We extend these theorems to a much more general class of sequences of random variables obeying a "threshold F^{α}-scheme" in which the distribution functions of the variables are close to the powers of a common F only in their right tails, above certain non-random non-decreasing threshold levels. Of independent interest is the characterization of the growth rate for extremal processes that we derived in order to be able to verify the conditions of our main theorem. We also establish the asymptotic pair-wise independence of record indicators in a special case of threshold F^{α}-schemes.

Online Optimal Pricing And Capacity Sizing For A $G/G/1$ Queue With Demand Learning

Y. Liu, North Carolina State University, NC, yliu48@ncsu.edu
X. Chen, Institute for Data and Decision Analytics, The Chinese University of Hong Kong, Shenzhen, chenxinyun@cuhk.edu.cn

We consider a pricing and capacity sizing problem in a $G/G/1$ queue with an unknown demand curve. The service provider's objective is to adaptively adjust the service price p and service rate μ so as to maximize cumulative expected revenues (the service revenue minus the delay penalty) over a given finite time horizon; in doing so, the service provider needs to resolve the tension between learning the unknown demand curve $\lambda(p)$ and maximizing earned revenues. In fact, such a problem has no analytic solution even when $\lambda(p)$ is unknown. In this talk, we develop an adaptive algorithm based on stochastic gradient descent (SGD) and provide asymptotic bounds for the regret (gap between the revenue when demand curve is known and the revenue earned using our online learning algorithm). We also discuss possible accelerating methods for the algorithm.

Heavy-Traffic Limits for Stationary Network Flows

W. Whitt, Columbia University, USA, ww2040@columbia.edu
W. You, Columbia University, USA, wy2225@columbia.edu

We establish heavy-traffic limits for the stationary flows in generalized Jackson networks, allowing an arbitrary subset of the queues to be critically loaded. The flows are the processes counting customers flowing from one queue to another or out of the network. The heavy-traffic limit with a single bottleneck queue is especially tractable because it yields limit processes involving one-dimensional reflected Brownian motion. That limit leads to accurate approximation of the index of dispersion for counts, which plays a crucial role in our new robust queueing network analyzer for approximating the steady-state performance of a non-Markovian open queueing network.

Nonstandard critical behavior of queue-based CSMA algorithms on the complete interference graph

E. Castiel, ISAE Supaero, IMT, France, eyal.castiel@isae.fr
S. Borst, Eindhoven University of Technology, Nokia Bell Labs, The Netherlands , s.c.borst@tue.nl
L. Miclo, CNRS, Toulouse School of Economics, France, Laurent.miclo@math.univ-toulouse.fr
F. Simatos, ISAE Supaero, France, florian.simatos@isae-supaero.fr
P. Whiting, Macquarie University, Australia philip.whiting@mq.edu.au

In this talk, we study an adaptive queue-based CSMA scheduling algorithm on the complete interference graph, which can be seen as a polling system with non-zero switchover time. The idleness induced by the decentralization of decisions entails non-standard heavy traffic behavior: when queue lengths are of order N we need to scale time by N^{a+1} for some $a \in (0,1)$ and the limit is deterministic. In contrast, the standard critical behavior corresponds to a time scale N^2 and the limiting process is typically a reflected Brownian motion. To prove this result we develop a new method to establish a stochastic averaging principle: the schedule evolves "much faster" than queue lengths and the proportion of the time a schedule is chosen is close (in an averaged way) to the invariant measure of the dynamic of the schedule (which depends on queue lengths). This method uses tools from functional analysis, in particular the log-Sobolev inequality. We then give a closed formula for a heavy traffic approximation by proving a state space collapse for the queue lengths (they live in fact close to a a one-dimensional manifold). We will present the procedure on a complete interference graph with critical arrival rates and deactivation rates polynomial in the queue length. We may also outline the challenges to overcome for a generalization of this result to different graphs.

Session 2.5 in Room M1 – Learning and Bandits
Chair: Sandeep Juneja

On The Futility of Dynamics In Robust Mechanism Design

S. Balseiro, Columbia Univeristy, srb2155@columbia.edu
A. Kim, Columbia University, aek2185@columbia.edu
D. Russo, Columbia, University, djr2174@gsb.columbia.edu

We consider a broad class of problems in which a principal repeatedly interacts with an agent who holds private information. The formulation includes as notable special cases (1) the problem of selling access to a digital good dynamically across several periods to a single buyer with unknown demand distribution and (2) the problem of dynamically designing labor contracts with an agent whose costs for producing output are unknown. Using the language of mechanism design, we allow the principal to design arbitrarily complex selling or contracting mechanisms. However, the agent will respond strategically to any dynamic mechanism in order to maximize their own benefit. Our main contribution is to formalize a broad impossibility result. If the principal's performance is measured of worst-case regret or a certain worst-case multiplicative ratio, the optimal mechanism for a T period problem simply repeats the optimal mechanism for single round problem. In this sense, the principal cannot benefit by using a mechanism with any meaningful dynamics, including any schemes that attempt to infer the private information of the agent.

Iterative Collaborative Filtering for Sparse Noisy Tensor Estimation

D. Shah, Massachusetts Institute of Technology, Cambridge, MA, devavrat@mit.edu
C. L. Yu, Cornell University, Ithaca, NY, cleeyu@cornell.edu

We consider the task of tensor estimation, i.e. estimating a low-rank 3-order $n \times n \times n$ tensor from noisy observations of randomly chosen entries in the sparse regime. In the context of matrix (2-order tensor) estimation, a variety of algorithms have been proposed and analyzed in the literature including the popular collaborative filtering algorithm that is extremely well utilized in practice. However, in the context of tensor estimation, there is limited progress. No natural extensions of collaborative filtering are known beyond “flattening” the tensor into a matrix and applying standard collaborative filtering. as long as each entry is observed independently with probability $p = \Omega(n^{-3/2+\epsilon})$ for any arbitrarily small $\epsilon > 0$. It turns out that $p = \Omega(n^{-3/2})$ is the conjectured lower bound as well as “connectivity threshold” of graph considered to compute similarity in our algorithm.

Learning to detect an oddball target

R. Sundaresan, Indian Institute of Science, Bangalore, India, rajeshs@iisc.ac.in

What policy should you employ if you were to locate an oddball target among many (two or more) distracters? You do not know the statistics of the observations that a ‘view’ of the oddball produces. Nor do you know the statistics of the observations arising from a view of a distracter. The objective is to minimise the search time, subject to keeping the probability of false detection low. Clearly, the common statistics of the distracter-generated observations should be compared, recognised as similar, and further recognised as different from those generated by the oddball target. But where should you look next, when should you stop, and when you stop, what should you decide as the oddball location? When the observations are points of a Poisson process (as in firings of neurons), and when the probability of false detection is driven down to zero, we will discuss the best growth rate of the expected search time. We will also identify an asymptotically optimal search policy that achieves the best growth rate. Finally, we will apply our results to a particular visual search experiment studied recently by neuroscientists and will highlight how prior information affects search performance.
This talk is based on joint work with N. K. Vaidhiyan.

Best arm selection, and regret minimization, for general distributions, in the bandit framework

S. Agarwal, Tata Institute of Fundamental Research, Mumbai, shubhada.agrawal@tifr.res.in
S. Juneja, Tata Institute of Fundamental Research, Mumbai, juneja@tifr.res.in

We consider the problems of the best arm selection under δ-correct guarantees, as well as regret minimization, in the multi-armed bandit framework. For both these class of problems, the lower bounds on sample complexity have been developed in the literature when the arms are allowed to have general distributions. However, matching asymptotic upper bounds exist only when the arm distributions belong to a somewhat restricted class of distributions such as the single parameter exponential family. In this talk, we relax this restriction to develop efficient algorithms when the underlying distributions have bounded $1 + \epsilon$ ($\epsilon > 0$) moments.

Session 2.6 in Room M2 – Approximating Rare Event Probabilities
Chair: Offer Kella

Counting Candy Crush Configurations

A. Hamilton, The University of Adelaide, Australia, a1668753@student.adelaide.edu.au
G. Nguyen, The University of Adelaide, giang.nguyen@adelaide.edu.au
M. Roughan, The University of Adelaide, Australia, matthew.roughan@adelaide.edu.au

In this talk we focus on the problem of counting the number of positions that can arise organically in the mobile phone game Candy Crush. This counting problem is called #CandyCrush. Due to the rare event nature of this problem Naïve Monte Carlo simulation will be ineffective unless an unreasonably large number of samples are used. In order to obtain accurate solutions in a reasonable amount of time, we design a Fully Polynomial Randomised Approximation Scheme (FPRAS) based on the Moser–Tardos algorithm and the Multilevel Splitting algorithm. This FPRAS allows us to approximate the solution to #CandyCrush to a desired level of accuracy in randomised polynomial time.

Sampling Conditional on a Rare Event

Z. Botev, UNSW Sydney, Australia, botev@unsw.edu.au

Suppose that we have at our disposal a well-designed importance sampling estimator of the probability of a rare event. We show how one can use such an importance sampling scheme to approximately simulate from the probability law, conditioned on the rare event (the zero-variance law). Our sampler creates an artificial regenerative process, whose limiting distribution is the zero-variance probability law. The rate of convergence of this sampler is determined by the large-deviations properties of the corresponding rare-event probability estimator. In analogy to the large-deviations properties of estimators (logarithmic efficiency; bounded relative error, and vanishing relative error), we define the corresponding large-deviations properties of the proposed rare-event sampler.

Rare Event Simulation for Steady-State Probabilities via Recurrency Cycles

K. Bisewski, CWI Amsterdam, The Netherlands, K.L.Bisewski@cwi.nl
D. Crommelin, CWI Amsterdam, The Netherlands, Daan.Crommelin@cwi.nl
M. Mandjes, University of Amsterdam, The Netherlands, M.R.H.Mandjes@uva.nl

We develop a new algorithm for the estimation of rare event probabilities associated with the steady-state (stationary) distribution of a Markov stochastic process with continuous state space $\mathbb{R}^d$ and discrete time steps (i.e. a discrete-time $\mathbb{R}^d$-valued Markov chain), e.g. numerical solution to Stochastic Differential Equation. The algorithm, which we coin Recurrent Multilevel Splitting (RMS), relies on the Markov chain's underlying *recurrent* structure (a concept akin to *regeneration*), in combination with the Multilevel Splitting method for rare event simulation. The algorithm is simple to use in practice and does not require detailed knowledge of the stochastic process, so that it can be applied to a broad class of systems (including 'black box' models that can be simulated numerically but that are too complex to be studied analytically). Extensive simulation experiments are performed for a 4-dimensional, nonlinear stochastic model that has some characteristics of complex climate models. The results show that RMS can boost the computational efficiency by several orders of magnitude compared to the standard Monte Carlo method.

Minimization of a stochastic convex function of a random variable subject to certain stochastic constraints

O. Kella, The Hebrew University of Jerusalem, Israel, Offer.Kella@huji.ac.il
R. Jacobovic, The Hebrew University of Jerusalem, Israel, royi.jacobovic@mail.huji.ac.il

We obtain a general solution to a problem of minimizing an integral of a nondecreasing stochastic process from zero to some nonnegative random variable S under the constraints that for some nonnegative random variable T, S is smaller than or equal to T almost surely and for some value t in $[0, ET]$ the expected value of S is t. The nondecreasing process, S and T are allowed to be dependent. Some special cases are considered and some examples are given. In particular, optimization related to the choice of clearing times in a clearing process and the optimization of certain quadratic functions with random coefficients subject so such constraints.

Session 2.7 in Rooms M5 and M6 – Asymptotics and Control
Chair: Asaf Cohen

Social Learning with Bandits

A. Sankararaman, University of Texas, Austin, abishek@utexas.edu
A. Ganesh, University of Bristol, UK, a.ganesh@bristol.ac.uk
S. Shakkottai, University of Texas, Austin, shakkott@austin.utexas.edu

We consider a multi-armed bandit problem where the number of arms, k, is large. There is also a large population of n agents collectively engaged in playing these arms. It is well known that the regret of a single agent over a time-horizon T scales as $k \log T$. Hence, if agents play individually with no communication, their collective regret scales as $kn \log T$. At the other extreme, with perfect information pooling, there is effectively a single agent, whose total regret scales as $k \log T$. Our main contribution is to show that a simple distributed scheme of recommendations achieves collective regret scaling as $k \log T$ polylog n.

Enhanced balancing of bias-variance tradeoff in stochastic simulation

H. Lam, Columbia University, USA, henry.lam@columbia.edu
X. Zhang, Columbia University, USA, zhang.xinyu@columbia.edu
X. Zhang, University of Science and Technology of China, China, zxh1998@mail.ustc.edu.cn

Biased stochastic estimators, such as finite-differences for noisy gradient estimation, often contain tuning parameters that balance impacts from the bias and the variance. While the optimal order of these parameters in terms of the simulation budget can be readily established, the precise best values depend on model characteristics that are typically unknown in advance. We investigate a framework to construct new estimators by combining simulation runs on sequences of tuning parameter values, such that the estimators consistently outperform a given tuning parameter choice in the conventional approach, regardless of the unknown model characteristics. We argue this outperformance via a minimax risk ratio and demonstrate, under this framework, a consistently outperforming estimator constructed by a weighting scheme that is characterized by a sum of two components with distinct decay rates.

Sensitivity analysis of reflected diffusions

D. Lipshutz, Technion — Israel Institute of Technology, Israel, lipshutz@technion.ac.il
K. Ramanan, Brown University, USA, Kavita_Ramanan@brown.edu

Reflected diffusions (RDs) constrained to remain in convex polyhedral domains arise in a variety of contexts, including as “heavy traffic” limits of queueing networks and in the study of rank-based interacting particle models. Sensitivity analysis of such an RD with respect to its defining parameters is of interest from both theoretical and applied perspectives. In this talk we characterize pathwise derivatives of an RD in terms of solutions to a linear constrained stochastic differential equation whose coefficients, domain and directions of reflection depend on the state of the RD. We demonstrate how pathwise derivatives are useful in Monte Carlo methods to estimate sensitivities of an RD, and also in characterizing sensitivities of the stationary distribution of an RD.

On singular control problems, the time-stretching method, and the weak-M1 topology

A. Cohen, University of Haifa, Israel, shloshim@gmail.com

We consider a general class of singular control problems with state constraints. Budhiraja and Ross (2006) studied a relaxed version of this class of problems and established the existence of optimal controls using the so-called ‘time-stretching’ method and the J1-topology. We show that the weak-M1 topology is better suited for establishing existence, since by using it, one bypasses the need for time-transformations, without any additional effort. Furthermore, we reveal how the time-scaling feature in the definition of the weak-M1 distance embeds the time-stretching scheme. This case study suggests that one can benefit from working with the weak-M1 topology in other singular control frameworks, such as queueing control problems under heavy traffic.

Session 2.8 in Rooms M7 and M8 – Revenue, Service, and Inventory Management
Chair: Qiong Wang & Marty Reiman

Size Matters, So Does Duration: The Interplay between Offer Size and Offer Deadline

Z. Hu, National University of Singapore, Singapore, bizhuz@nus.edu.sg
W. Tang, National University of Singapore, Singapore, wenjie.tang@nus.edu.sg

We study a Stackelberg game involving a proposer and responder. The proposer acts first by making an offer to the responder with a deadline, and the responder, following a continuous time finite-horizon search for alternative offers, has to respond to the offer by the deadline. We find that the proposer should offer a non-exploding offer when the alternative offer market is favorable to the responder. When the alternative offer market is unfavorable to the responder, the proposer can profit from making an exploding offer with a smaller size in a harsher market. We further study the case where the responder has private knowledge of the alternative offers' arrival rate. A sign-up bonus type of contract can serve as a self-selection mechanism and we characterize the optimal contract.

Asymptotic optimality of a local scheduling policy for a queueing system with customer feedback

J. Huang, The Chinese University of Hong Kong, junfeih@cuhk.edu.hk

Consider a "W"-queueing systems with 3 customer classes, 2 servers and customer feedback. After the current service is finished, a customer may transfer to another customer class, or leave the system. The congestion is measured by the cumulative queueing costs. We assume that each server has access to the **local information** only in the sense that only its own queue lengths are available but not all the three. We propose a generalized $c\mu$-type policy and prove that it is asymptotically optimal among policies which may also use the information of all the queue lengths.

A Re-solving Heuristic with Uniformly Bounded Loss for Network Revenue Management

P. Bumpensanti, Georgia Institute of Technology, pornpawee@gatech.edu
H. Wang, Georgia Institute of Technology, he.wang@isye.gatech.edu

We consider a canonical (quantity-based) network revenue management problem, where a firm accepts or rejects incoming customer requests irrevocably in order to maximize expected revenue given limited resources. Due to the curse of dimensionality, the exact solution to this problem by dynamic programming is intractable when the number of resources is large. We study a family of re-solving heuristics that periodically re-optimize an approximation to the original problem known as the deterministic linear program (DLP), where random customer arrivals are replaced by their expectations. We find that, in general, frequently re-solving the DLP produces the same order of revenue loss as one would get without re-solving, which scales as the square root of the time horizon length and resource capacities. By re-solving the DLP at a few selected points in time and applying thresholds to the customer acceptance probabilities, we design a new re-solving heuristic whose revenue loss is uniformly bounded by a constant that is independent of the time horizon and resource capacities.

Wednesday 16:30 – 18:00

Session 3.1 in Rooms P1 and P2 – Efficiency in Stochastic Optimization and Simulation Experiments
Chair: Henry Lam

Efficient estimation of sensitivity of tail risk measures with applications in risk averse stochastic optimization

A. Deo, Tata Institute of Fundamental Research, Mumbai, anand.deo@tifr.res.in
K. Murthy, Singapore University of Technology and Design, Singapore, karthyek@gmail.com

We consider tail risk averse stochastic optimization formulations which aim to identify decisions achieving low average risk and low tail risk simultaneously. Such formulations are useful in diverse contexts ranging from mitigating tail risks in portfolio optimization to ensuring fairness for minority subpopulations in machine learning tasks. The problem is challenging in data-driven settings due to the lack of samples representing tail regions. Motivated by the tail extrapolation approaches that are prevalent in extreme value theory, we identify limiting distributions for suitably normalized versions of the objective (representing risk) and its gradient, when conditioned on the objective being unusually large. We then utilise the resulting approximations to efficiently estimate tail risk measures such as CVaR and its gradients in a data-driven fashion. Most of the existing works either begin assuming a stochastic model (or) estimate CVaR as a plugin from the empirical distribution. While the former approach largely ignores the question of how to arrive at a stochastic model that has high fidelity in the tails, the latter is ineffective in that the tail regions are usually not well represented in the data. The utilization of extreme value based extrapolation of gradients to execute a first-order descent method is entirely novel in the setting considered and it crucially bridges the gap between the two extreme approaches considered in the literature.

Space-filling design for non-linear models

C.-H. Rhee, Northwestern University, USA, chang-han.rhee@northwestern.edu
E. Zhou, Georgia Tech, USA, enlu.zhou@isye.gatech.edu
P. Qiu, Georgia Tech and Emory University, USA, peng.qiu@bme.gatech.edu

Traditional space-filling designs for computer experiments aim to fill the parameter space with design points that are as “uniform” as possible. However, the resulting design points may be non-uniform in the model output space failing to provide a reliable representation of the output manifold, and becoming highly inefficient or even misleading in case the computer experiments are non-linear. In this talk, we propose and analyze an iterative algorithm that fills in the model output manifold uniformly—rather than the parameter space uniformly—so that one could obtain a reliable understanding of the model behaviors with the minimal number of design points. This problem and the proposed solution has close connection to the rare-event simulation problems and non-parametric importance sampling algorithms. This connection leads to a novel convergence analysis of non-parametric importance sampling algorithms in Wasserstein metric.

Rare Event Estimation for Elliptical Random Vectors

D. Chen, Purdue University, West Lafayette IN, chen1407@purdue.edu
R. Pasupathy, Purdue University, West Lafayette IN, pasupath@purdue.edu
H. Honnappa, Purdue University, West Lafayette, honnappa@purdue.edu

Monte Carlo estimation of rare events in high-dimensional multivariate settings is complicated. In this talk, we consider the estimation of the likelihood that a random vector lies in an ‘atypical’ closed convex or concave set - the ostensible rare event of interest. Since standard Monte Carlo estimation is not a viable technique, we focus on importance sampling (IS) methods that exploit the (local) geometry of the atypical set to identify a dominating point, around which we can define a proposal distribution for importance sampling. We illustrate this method on several elliptical multivariate distributions.

Bounding Optimality Gaps via Bagging

H. Lam, Columbia University, New York, US, henry.lam@columbia.edu
H. Qian, Columbia University, New York, US, h.qian@columbia.edu

We investigate a statistical method to assess the suboptimality of a given solution in data-driven stochastic optimization, via estimating the solution's optimality gap. Our approach is based on bootstrap aggregating, or bagging, resampled sample average approximation (SAA). We show how this approach generates valid statistical confidence bounds for non-smooth optimization, and demonstrate and compare its statistical efficiency and stability with some existing methods. We also present our theory by viewing SAA as a kernel in an infinite-order symmetric statistic, which leads to some generalizations of classical central limit results for SAA.

Session 3.2 in Room P3 – Dynamic Investment and Stochastic Networks
Chair: Agostino Capponi

Dynamic Investment and Financing Decisions with Internal and External Liquidity Management

N. Chen, The Chinese University of Hong Kong, Hong Kong, nchen@se.cuhk.edu.hk
J. Ji, The Chinese University of Hong Kong, Hong Kong, jhji@se.cuhk.edu.hk
Y. Tian, Ryukoku University, Japan, tian@econ.ryukoku.ac.jp

We develop a theoretical model of dynamic investments, dividend payouts, debt borrowing, external equity financing, and bankruptcy for financially constrained firms. The model characterizes the central importance of liquidity management in corporate decision making in the presence of external financing costs. Mathematically, to solve for the recursive equilibrium of the problem, we formulate it as an optimal stopping problem with a fixed-point structure embedded.
Our model can generate rich implications. Particularly, we find that the debt may yield two opposing effects to the firm's investment decisions if it has limited liquidity. On one hand, debt issuance may enhance the size of current investment; on the other hand, debt may reduce the actual profit to the firm. Our model characterizes quantitatively how these two effects, interacting with the cash management, will shape up the firm's investment, financing, bankruptcy, and payout decisions. The paper also discusses the implications of liquidity and leverage requirements in the current banking regulatory framework.

A Dynamic Network Model of Interbank Lending – Systemic Risk and Liquidity Provisioning

A. Capponi, Columbia University, USA, ac3827@columbia.edu
X. Sun, Columbia University, USA, xs2235@columbia.edu
D. D. Yao, Columbia University, USA, yao@columbia.edu

We develop a dynamic model of interbank borrowing and lending activities in which banks are organized into clusters, and adjust their monetary reserve levels so as to meet prescribed capital requirements. Each bank has its own initial monetary reserve level and faces idiosyncratic risks characterized by an independent Brownian motion; whereas system wide, the banks form a hierarchical structure of clusters. We model the interbank transactional dynamics through a set of interacting measure-valued processes. Each individual process describes the intra-cluster borrowing/lending activities, and the interactions among the processes capture the inter-cluster financial transactions. We establish the weak limit of the interacting measure-valued processes as the number of banks in the system grows large. We then use the weak limit to develop asymptotic approximations of two proposed macro-measures, the liquidity stress index and the concentration index, both capturing the dynamics of systemic risk. We use numerical examples to illustrate the applications of the asymptotics and related sensitivity analysis with respect to various indicators of financial activity.

Disruption and Rerouting in Supply Chain Networks

J. R. Birge, Booth School of Business, University of Chicago, Chicago, IL 60637, USA, john.birge@chicagobooth.edu
A. Capponi, Industrial Engineering and Operations Research Department, Columbia University, New York, NY10027, USA, ac3827@columbia.edu
P.-C. Chen, Department of Industrial and Manufacturing Systems Engineering, The University of Hong Kong, Pokfulam, Hong Kong, pcchen@hku.hk

We study systemic risk in a supply-chain network where firms are connected through purchase orders. Each firm's short-term investment return is subject to a random shock leading to a default. The shock propagates through the supply-chain network via input-output linkages between buyers and suppliers. Firms endogenously take contingency plans to mitigate the impact generated from disruptions. They reroute undelivered orders to alternative buyers and switch excess demand to different suppliers. An equilibrium is reached when contagion from disruption stops. We develop an algorithm to recover the Pareto dominant equilibrium with the greatest amount of delivered orders. We show that lower concentration of orders results in a more fragile network if connectivity is high and firms are highly capitalized. If firms are lowly capitalized, however, lower concentration of orders reduces network fragility. Highly concentrated orders are always preferred in networks with low connectivity.

A concentration phenomenon for system failure times under simultaneous shocks

J. Barrera, Universidad Adolfo Ibáñez, Chile, javiera.barrera@uai.cl
G. Lagos, University of Santiago, Chile, guido.lagos@usach.cl

We consider the Lévy-frailty Marshall–Olkin distribution to model the times of occurrence of shocks that hit or affect agents of a system, where each shock can simultaneously affect one or more agents at a time. In this setting we define the system failure time as the time when the k-th out of a total of n agents is hit or shocked, where k is sufficiently close (or equal to) n. Our main contribution is a concentration result for the system failure time, as the time and size of the system grow together in a non-trivial scaling regime. Our result allows to study common-cause shocks or failures in systems in a relevant and practical fashion: it allows to estimate and give confidence intervals for the system failure time and at same time only needs to specify a reduced family of parameters.

Session 3.3 in Room P4 – Level-Dependent QBD Processes
Chair: Yonit Barron

Drift Conditions for the Ergodicity of a Class of Level-Dependent $M/G/1$-type Processes

J. Cordeiro, University of Dayton, U.S.A., jcordeiro1@udayton.edu
J. Kharoufeh, University of Pittsburgh, U.S.A., jkharouf@pitt.edu

Structured Markov chains, such as quasi-birth-and-death (QBD) processes, $M/G/1$-type processes, and $G/M/1$-type processes form the basis of matrix-analytic methods that are often used for the performance evaluation and stability of queueing systems. While closed-form drift conditions for the ergodicity of level-independent versions of these processes are known, the level-dependent case is not as clear. Recently, Cordeiro, et al. (2019) examined level-dependent QBDs (LDQBDs) by examining the Markov chain embedded at jump epochs in which block terms of the transition probability matrix converge over levels. In this talk, we will describe how a similar Foster-Lyapunov drift approach can be applied to establish closed-form ergodicity conditions for the level-dependent version of the $M/G/1$-type process. Lastly, several approaches to weakening the regularity condition for block convergence over levels of the process will be discussed, particularly for LDQBD processes.

Asymptotic Analysis of QBDs with Time-Varying Periodic Rates with Examples

B. Margolius, Cleveland State University, Cleveland, OH, USA, b.margolius@csuohio.edu

In this talk we outline an approach for obtaining asymptotic estimates of level probabilities of continuous time level independent Quasi-Birth-Death processes with time-varying periodic transition rates. We apply the approach to examples where an explicit formula for the corresponding random walk process is available. Among the examples we consider is a multiserver queue with vacations, in which the server works with different service rates rather than completely terminates service during his vacation period.
The approach applies techniques of Flajolet and Sedgwick outlined in *Analytic Combinatorics* and involves finding the roots of the determinant of a matrix related to the generating function for a two-dimensional random walk over a single time period.
The level probabilities as the level number tends to infinity tend to periodic functions of the form $\sum_j r_j^n \Pi_{k,j}(n) f_{j,k}(t)$ where n is the level number and $f_{j,k}(t)$ is a function of time within the period. $f_{j,k}(t)$ depends on the phase k and the singularity r_j. r_j and $\Pi_{j,k}(n)$ do not depend on time. The expressions $\Pi_{j,k}(n)$ are polynomials in n.

Extremal Queues Given First Two Moments

Y. Chen, Columbia University, USA, yc3107@columbia.edu
W. Whitt, Columbia University, USA, ww2040@columbia.edu

This paper studies upper bounds for the mean (steady-state and transient) waiting time in the GI/GI/1 queue given the first two moments of the interarrival-time and service-time distributions. For distributions with support on bounded intervals, we show that the upper bounds (with one distribution given and overall) are attained at distributions with support on at most three points. The proof exploits fixed point theory and optimization theory in addition to standard stochastic theory for the model. We then apply relatively tractable numerical algorithms to identify the optimal distributions within that class. For the overall upper bound with unbounded support sets, we propose a simple approximation formula and provide a numerical comparison of the approximations and bounds, showing that the new approximate bound is very accurate.

A threshold policy in a Markov-modulated production system with server vacation: The case of continuous and batch supplies

Y. Barron, Ariel University, Ariel 40700, Israel, ybarron@ariel.ac.il

We consider a Markov-modulated fluid flow production model under D-policy, that is, as soon as the storage reaches level 0, the machine becomes idle until the total storage exceeds a predetermined threshold D. Thus, the production process alternates between a busy and an idle machine. During the busy period the storage decreases linearly due to continuous production and increases due to supply; during the idle period no production is rendered by the machine and the storage level increases only by supply arrivals. We consider two types of models with different supply process patterns: continuous inflows with linear rates (fluid type), and batch inflows, where the supplies arrive according to a Markov additive process and their sizes are independent and have phase-type distributions depending on the type of arrival (MAP-type). Four types of costs are considered: a setup cost, a production cost, a penalty cost for an idle machine, and a cost for the storage. Based on tools from multi-dimensional martingales and hitting times theory, we derive explicit formulas for these cost functionals in the discounted case. Numerical examples, a sensitivity analysis, and insights are provided.

Session 3.4 in Room P5 – Markov Chains
Chair: Ivo Adan

Overview of first passage times calculations for double thresholded processes

S. Kapodistria, Eindhoven University of Technology, The Netherlands, s.kapodistria@tue.nl

In this presentation, we analyse a multiplicative growth-collapse process X that grows linearly in time and experiences downward jumps at Poisson epochs. The collapses are modeled by multiplying the present process position by a fixed proportion $Q \in (0, 1)$. For this process and its appropriate reflected version, we solve exit problems that concern identification of laws of exit times from fixed intervals. We also consider the first downward passage time when the process starts from its stationary distribution.
This is joint work with Remco van der Hofstad, Zbigniew Palmowski, and Seva Shneer.

Algebraic Classification of Markov Chains for Assessing Quasi-Birth–Death Process Tractability

A. Wickeham, University of Minnesota, U.S., wicke086@umn.edu
S. Doroudi, University of Minnesota, U.S., sdoroudi@umn.edu

Infinite state repeating Markov chains, particularly Quasi-Birth–Death processes, are frequently used in the modeling of service systems. In the majority of the literature studying such processes, the chains' limiting probability distributions and related metrics of interest are determined via numerical methods. A smaller subset of the literature has provided methods for determining such quantities exactly for chains exhibiting specific structures. While numerical methods are generally more broadly applicable, exact analysis sometimes allows for obtaining closed-form solutions; such solutions can often provide greater insight into the chains' structural properties.
Our work in progress observes that many such exact solution techniques are in principle applicable to a broader class of chains (than originally designed for), but only through solving a potentially intractable system of higher order polynomials. We introduce a paradigm that classifies Markov chains by their algebraic properties, with the goal of determining for which chains exact analysis is tractable. Specifically, we seek to characterize the structure of the polynomial system that must solved in order to obtain a solution, regardless of the exact method of choice. Our classification scheme could provide practitioners with a simple way of determining whether a chain of interest is suitable for exact analysis.

Stability condition for a two class discrete-time queueing model with bounded variable-length service times

S. Dendievel, Ghent University, Belgium, Sarah.Dendievel@UGent.be
H. Bruneel, Ghent University, Belgium, Herwig.Bruneel@UGent.be

We consider a discrete-time queueing model with one single infinite waiting room and two servers. There are two classes of customers, each having their own dedicated server. We assume that the types of two consecutive customers are interclass-correlated and that the customers are served in their order of arrival. The impact of non-independent arrivals may cause a blocking effect: a class clustering in the arrivals can block customers of the other class arriving later in the waiting room regardless of the state of the other class' server. Our aim in the present work is to formulate a stability criterion for a two-class queueing system when the service times are of variable-length and bounded by $s \geq 1$ time slots. We prove the impact of class clustering in the arrival stream on the stability of the system and illustrate our results with numerical examples.
In the recent literature on discrete-time multi-class queueing models, deterministic service times and geometric service times have been considered. For deterministic service times, a Markov state description is introduced to keep track of the state of the server and explicit expressions for the probability generating functions (pgfs) of performance measures are derived. For geometric service times, the randomness of the service times increases the difficulty of the analysis compared to deterministic case. For instance, customers may leave the system after being served in an order different of their arrival order. This creates the need of additional information on the system state description. The state description for the geometric case is determined by a vector of three variables that leads to explicit expressions of the pgfs for the system characteristics. However, the analysis developed for the geometric case does not seem to be suitable for the case of bounded variable-length service times. In this paper, we obtain our results by defining an appropriate state representation of the service times of the two servers and by using a matrix-analytic approach.

Optimizing Activity Sequences by Minimizing Rework Effects

I. Adan, Eindhoven University of Technology, Netherlands, I.Adan@tue.nl
T. Wilschut, Eindhoven University of Technology, Netherlands, T.Wilschut@tue.nl

Large scale design and construction projects are often faced with major schedule and budget overruns. One of the causes of these overruns is rework. Rework occurs when an activity is executed incorrectly the first time and (partially) has to be redone. Rework is caused by the dependencies between different activities in a schedule. This network of dependencies is called the process architecture. To visualise these dependencies the Design Structure Matrix is used. This talk focusses on optimizing an activity sequence by minimizing rework effects. To efficiently calculate the effect of an activity sequence on the duration and cost of a project, a Markov chain model is defined. This is done by using the error probabilities of information transfers between activities. The sequence of activities is optimized for different objective functions including mean and variance of the total project duration.

Session 3.5 in Room M1 – Statistical and Computational Aspects of Modern Machine Learning Problems
Chair: Fred Roosta-Khorasani

Communication-Efficient Distributed Optimization Methods for Large-Scale Machine Learning

R. Crane, University of Queensland, Australia, r.crane@uq.edu.au
F. Roosta, University of Queensland, Australia, fred.roosta@uq.edu.au

Continuous technological and communication advancements have enabled the collection of and access to ever-growing large scale datasets. There is a significant amount of research and development being devoted to machine learning problems in general, and the underlying optimization algorithms in particular, that are formulated on these large scale datasets. However, lack of adequate computational resources, in particular storage, can severely limit, or even prevent, any attempts at solving such optimization problems in a traditional stand-alone way, e.g., using a single machine. This can be remedied through distributed computing, in which resources across a network of stand-alone computational nodes are "pooled" together so as to scale to the problem at hand. Optimization algorithms designed for distributed environments can efficiently leverage the computational resources of a network of machines. A significant issue with this framework of computing is that inter-machine communication can be expensive. Therefore, in such environments it is necessary to design algorithms that use a low number of communication rounds. In this talk, I will discuss the significant potential of communication-efficient distributed second-order optimization methods and discuss my research in developing such methods that remedy many of the shortcomings of the existing methods.

Implicit Langevin algorithms for sampling from log-concave densities

L. Hodgkinson, The University of Queensland, Australia, liam.hodgkinson@uqconnect.edu.au
R. Salomone, The University of New South Wales, Australia, r.salomone@unsw.edu.au
F. Roosta, The University of Queensland, Australia, fred.roosta@uq.edu.au

Effectively sampling from unnormalised probability distributions is a fundamental aspect of the Monte-Carlo method, and is central in Bayesian inference. We study sampling schemes derived from the implicit θ-method discretisation method applied to the overdamped Langevin diffusion stochastic differential equation, in the context of log-concave target densities. Such densities commonly arise in Bayesian regression problems involving generalised linear models, such as Bayesian logistic or softmax regression. Obtaining subsequent samples amounts to solving a strongly-convex optimisation problem, which is readily achievable using one of numerous existing methods. These implicit schemes enjoy several theoretical advantages over the competition; in particular, the unadjusted Langevin algorithm, which has seen a recent surge of popularity for high-dimensional sampling problems. Such advantages include geometric ergodicity and stability when $\theta \geq 1/2$ for *all* step sizes. Numerical examples supporting our theoretical analysis will also be provided.

Probabilistic Symmetries in Neural Networks

R. Tsuchida, The University of Queensland, Australia, s.tsuchida@uq.edu.au
F. Roosta, The University of Queensland, Australia, fred.roosta@uq.edu.au
M. Gallagher, The University of Queensland, Australia, marcusg@itee.uq.edu.au

The multilayer perceptron (MLP) is a widely used machine learning model where an input vector is transformed to an output through a series of linear and non-linear mappings. The linear mappings are represented by matrices, called the weights, and the non-linear mappings are called activations.
A typical way in which an MLP is used is to randomly initialise the weights from some distribution and optimise the weights using stochastic gradient descent according to some regularised objective on the output. The class of functions representable at the output is determined by the inner product in the hidden layers of the MLP, known as the kernel.
Another theoretical construction analyses the distribution over functions induced by placing an IID distribution over the weights. In the appropriate limit as the number of rows of the first matrix goes to infinity, the resulting distribution over functions has the interpretation of a prior distribution of a Gaussian process with a specific kernel.
Motivated by these two use cases, we analyse the kernel in an appropriate limit when the weights follow different probabilistic symmetries: isotropic Gaussian, rotatable, IID, and exchangeable. We show that the weights of MLPs trained using stochastic gradient descent follow a a type of exchangeability.

A new fast algorithm to approximate the leverage scores of big time series data: theory and application

A. Eshragh, The University of Newcastle, Australia, ali.eshragh@newcastle.edu.au
F. Roosta, The University of Queensland, Australia, fred.roosta@uq.edu.au
M. Mahoney, The University of California at Berkeley, USA, mmahoney@stat.berkeley.edu
A. Nazari, Deakin University, Australia, asef.nazari@deakin.edu.au

We utilize randomized numerical linear algebra techniques to develop a new fast algorithm to compute the leverage scores of an autoregressive model approximately. We show that the time complexity of this new algorithm is $\mathcal{O}(1)$ and the accuracy of approximations lies within $(1+\varepsilon)$ of the true leverage scores with high probability. These theoretical results are exploited to fit autoregressive models to big time series data by subsampling the data effectively and reducing its dimensionality significantly. Empirical results on synthetic as well as real world data highly support the theoretical results and reveal the efficacy of this new approach.

Session 3.6 in Room M2 – MCMC Algorithms
Chair: Michael Choi

Constrained Hamiltonian Monte Carlo for PDE Inverse Problems

K. X. Au, National University of Singapore, Singapore, khai@u.nus.edu
M. M. Graham, National University of Singapore, Singapore, m.m.graham@nus.edu.sg
A. H. Thiery, National University of Singapore, Singapore, a.h.thiery@nus.edu.sg

We consider the problem of inference in models of the form $y = \mathcal{G}(x) + \eta$ where y is a vector of observed variables, x is the quantity of interest to be inferred, $\mathcal{G}$ is the observation operator and $\eta \sim \mathcal{N}(0, \sigma^2 I)$ is a vector of additive noise in the observations. Specifically, we consider the case where $\mathcal{G}$ involves solving a system of partial differential equations (PDEs) in x. In such models, the posterior mass concentrates in the neighbourhood of a manifold as σ tends to zero. As a result, the efficiency of MCMC methods deteriorates due to the need to take increasingly small steps.
In this work, we present a constrained HMC algorithm that is robust to small σ values, i.e. low noise. Taking the observations generated by the model to be constraints on the prior, we define a manifold on which the constrained HMC algorithm generate samples. By exploiting the geometry of the manifold, our algorithm is able to take larger step sizes than more standard MCMC methods, resulting in a more efficient sampler. We explain the setup of manifold HMC and additional steps necessary to ensure the reversibility of the resulting Markov chains. Finally, we apply the algorithm to several PDE inverse problems to verify its efficiency in the low noise regime.

Importance sampling for partially observed Markov chain models

A. J. Black, University of Adelaide, Australia, andrew.black@adealide.edu.au

Particle marginal Metropolis Hastings (pmMH) is a popular method for performing Bayesian inference when it is only possible to simulate from the underlying model. This uses a particle filter to estimate the likelihood, which is then used in a basic Metropolis Hastings algorithm. A well know problem with particle filters is that the more accurate the observations of a system, the worse the filter performs. For models where we observe a single component of the state exactly, the cost of producing simulations becomes very high. In this talk I will discuss some recent work using importance sampling to generate realisations of Markovian models that match observations exactly. When used in a pmMH scheme, the importance sampling provides a large speed-up in terms of the effective sample size per unit of computational time compared to simple bootstrap sampling.

Hitting time, mixing time and geometric interpretations of Metropolis-Hastings reversiblizations

M. Choi, Institute for Data and Decision Analytics, The Chinese University of Hong Kong, Shenzhen, China, michaelchoi@cuhk.edu.cn
L. Huang, Fujian Normal University, Fuzhou, Fujian, China, huanglj@fjnu.edu.cn

Given a target distribution μ and a proposal chain with generator Q on a finite state space, in this talk we study two types of Metropolis-Hastings (MH) generator $M_1(Q,\mu)$ and $M_2(Q,\mu)$ in a continuous-time setting. While M_1 is the classical MH generator, we define a new generator M_2 that captures the opposite movement of M_1 and provide a comprehensive suite of comparison results ranging from hitting time and mixing time to asymptotic variance, large deviations and capacity, which demonstrate that M_2 enjoys superior mixing properties than M_1. To see that M_1 and M_2 are natural transformations, we offer an interesting geometric interpretation of M_1, M_2 and their convex combinations as ℓ^1 minimizers between Q and the set of μ-reversible generators, extending the results by Billera and Diaconis (2001). We provide two examples as illustrations. In the first one we give explicit spectral analysis of M_1 and M_2 for Metropolised independent sampling, while in the second example we prove a Laplace transform order of the fastest strong stationary time between birth-death M_1 and M_2.

Equity Valuation and Falling Cost of Data Analytics

J. Keppo, National University of Singapore, Singapore, keppo@nus.edu.sg
H. M. Tan, National University of Singapore, Singapore, thm@u.nus.edu

We model two risk-averse decision makers who optimize the size of their investment, and the level of information on the investment through costly data analytics. The decision makers affect the market price of the investment, and optimize their own expected utility given the other decision maker's quantity of data analytics. We study the resulting market equilibrium. We show that in this game model, the falling cost of data analytics raises the Nash equilibrium quantity of data analytics for each decision maker. This results in an increase in the expected equilibrium price of the investment. Hence, our model gives a new explanation for the rise of equity valuation during the last few decades - the falling cost of data analytics. We show empirical evidence for this by using U.S. equity prices.

Session 3.7 in Rooms M5 and M6 – Control of Stochastic Systems
Chair: Hayriye Ayhan

Scheduling Same-day Deliveries Using Truck-to-Drone Re-supply

D. F. Silva, Auburn University, AL, USA, silva@auburn.edu
A. Smith, Auburn University, AL, USA, smithae@auburn.edu

Consider an urban order-delivery operation with a single depot. Customers, at random locations within the city, place orders at random times. Each day, the depot works for a fixed number of hours. If any orders arrive outside operating hours, they are available at the beginning of the next delivery period. We assume that there is a reward for each delivery completed. Furthermore, orders that are not delivered by the end of the delivery period are discarded and a penalty is incurred. There is a single truck available and the orders must be delivered by the truck. Additionally, there is a heavy-duty drone that can take orders from the depot to the truck. We assume that travel times between customers and the depot are deterministic, but can be different for the truck and the drone. We also assume that the drone can only rendezvous with the truck at certain locations. We model this system as a Markov Decision Process. We use the MDP model to find optimal policies for small instances of the problem. We also develop heuristic approaches for large instances.

A Little Redundancy Goes A Long Way: Convexity in Redundancy Systems

K. Gardner, Amherst College, kgardner@amherst.edu
E. Hyytiä, University of Iceland, esa@hi.is
R. Righter, University of California, Berkeley, rrighter@berkeley.edu

Redundancy is an increasingly popular technique for reducing response times in computer systems, and there is a growing body of theoretical work seeking to analyze performance in systems with redundancy. The idea is to dispatch a job to multiple servers at the same time and wait for the first copy to complete service. Redundancy can help reduce response time because redundant jobs get to experience the shortest of multiple queueing times and potentially of multiple service times—but it can hurt jobs that are not redundant and must wait behind the redundant jobs' extra copies. Thus in designing redundancy systems it is critical to find ways to leverage the potential benefits without incurring the potential costs. Scheduling represents one tool for maximizing the benefits of redundancy. We study a scheduling policy called Least Redundant First (LRF), under which less-redundant jobs have priority over more-redundant jobs. Our goal is to understand the marginal impact of redundancy: how much redundancy is needed to get the biggest benefit? We provide a surprisingly intricate proof that mean response time is convex as well as decreasing as the proportion of jobs that are redundant increases under LRF for exponential services. On the other hand, the same is not true under first-come first-served scheduling: in this case, mean response time may be neither decreasing nor convex, depending on the parameter values. Thus, the scheduling policy is key in determining both whether redundancy helps and the marginal effects of adding more redundancy to the system.

Analysis of Batch Dispatch Policies for Order Fulfillment

N. Gautam, Texas A&M University, USA, gautam@tamu.edu

In this research we consider orders that arrive one by one over time to a fulfilment center. Each order is associated with a product that needs to be delivered expeditiously to nearby locations using vehicles. However, each vehicle could batch multiple orders and deliver them in a single trip. Benefits of batching include higher capacity, fewer vehicles, and lower environmental impact. The main drawbacks of batching is the reduced quality of service due to delay in delivery. To address this trade off, our objective is to formulate an optimal control problem and develop a policy for batching and dispatching that would minimize the long-run average cost per unit time. For the optimal policy, our state variable not only includes the number of orders but also incorporates when they arrived, and the delivery locations. We model the stochastic dynamics of the state using partial differential equations. Using that we evaluate the performance of a threshold policy and obtain the threshold that minimizes the long-run average cost per unit time which we compute using a renewal reward approach. We also evaluate various delivery mechanisms such as first come first served and shortest Hamiltonian path. In addition, we evaluate the effectiveness of including all the information for our decisions, as opposed to just the number in the system.

Realtime Integrated Learning and Decision Making for Deteriorating Systems

C. Drent, Eindhoven University of Technology, The Netherlands, c.drent@tue.nl
M. Drent, University of Luxembourg, Luxembourg, melvin.drent@uni.lu
J. Arts, University of Luxembourg, Luxembourg, joachim.arts@uni.lu
S. Kapodistria, Eindhoven University of Technology, The Netherlands, s.kapodistria@tue.nl

We focus on the decision making for condition based maintenance of deteriorating single-component systems. Such systems although seemingly identical, they are oftentimes heterogeneous due to different exogenous (e.g., environment) and endogenous (e.g., quality) conditions. This is captured in their corresponding deterioration (i.e. degradation) process rendering the underline model observable but unknown (uncertain). To this end, it is paramount to integrate the learning of the individual degradation model with the maintenance optimization, and to strive for tailored maintenance decisions on an individual system level. This is in sharp contrast with the conventional approach, appearing in the maintenance literature, of separately (typically sequentially) modeling the degradation patterns and performing the maintenance optimization. It is notable, that under this conventional approach, it is implicitly assumed that each component has the same a-priori degradation process, based on which maintenance decisions are subsequently made for the whole population.
In this talk, as a vehicle of illustration of the underlying ideas and the approach, we assume that components degrade according to a compound Poisson process with geometric compounding distribution. The parameters of the degradation process, for which we only have information in the form of prior distributions, differ for each component and are inferred with increasing accuracy using real-time degradation data in a Bayesian framework. This is incorporated in the decision making problem in which the objective is to minimize the total expected discounted (corrective and preventive) replacement costs over an infinite horizon. Using the theory of Markov decision processes, by appropriately incorporating the population heterogeneity, we show that the optimal replacement policy for this class of problems is a monotonically non-decreasing control limit policy that depends on the individual degradation trajectory of a component. In an extensive numerical experiment, we benchmark our model against straightforward approaches that either neglect population heterogeneity or are naive in the way that parameters are estimated. Finally, we highlight the practical value of the optimal replacement policy by applying the model to real-life degradation data of a component crucial for the operation of medical imaging systems.

Session 3.8 in Rooms M7 and M8 – Applied Probability in Financial Engineering
Chair: Lingfei Li

A multi-factor regime switching model for inter-trade durations in the limit order market

X. Chen, Chinese University of Hong Kong, Shenzhen, China, chenxinyun@cuhk.edu.cn
Z. Li, Hunan University, China, lzcj18@aliyun.com
H. Xing, Stony Brook University, US, xing@ams.sunysb.edu

We identify a common bimodal distribution for inter-trade durations in recent data from the Nasdaq limit order market. This novel finding has led us to build a parsimonious regime switching model in which the inter-trade durations move between long- and short-duration regimes. Such regime switching is probably caused by the response of trading algorithms to market information. In order to analyze which types of information in the limit order market do trading algorithms learn from and react to, we incorporate limit order book (LOB) information as factors into the model by setting the transition matrix to be time-varying and factors dependent. Empirical study for the Nasdaq stocks suggest that the model has good in-sample fitness and the out-of-sample performance is better than benchmark duration models. Finally, our analysis on the significances of LOB factors contribute to the study of market microstructure in a high frequency world.

Langevin Algorithms and Momentum-Based Acceleration for Non-convex Stochastic Optimisation

X. Gao, Chinese University of Hong Kong, Hong Kong, xfgao@se.cuhk.edu.hk
M. Gurbuzbalaban, Rutgers University, USA, mg1366@rutgers.edu
L. Zhu, Florida Sate University, USA, zhu@math.fsu.edu

We consider stochastic non-convex optimization problems that arise in several applications including machine learning and the stochastic gradient Hamiltonian Monte Carlo (SGHMC) algorithm to solve them. We obtain the first finite-time global convergence guarantees for SGHMC in the context of both empirical and population risk minimization. Our results show SGHMC can achieve acceleration on a class of non-convex problems compared to overdamped Langevin MCMC approaches such as the stochastic gradient Langevin dynamics.

Analysis of Markov Chain Approximation for Diffusion Models with Non-Smooth Coefficients

G. Zhang, The Chinese University of Hong Kong, Shenzhen, zhanggongqiu@cuhk.edu.cn
L. Li, The Chinese University of Hong Kong, Hong Kong, lfli@se.cuhk.edu.hk

Diffusion models with non-smooth coefficients often appear in financial applications, with examples including but not limited to threshold models for financial variables, the pricing of occupation time derivatives and shadow rate models for interest rate dynamics. To calculate the expected value of a discounted payoff under general state-dependent discounting and monitoring of barrier crossing, continuous time Markov chain (CTMC) approximation can be applied. In a recent work, Zhang and Li (2018, forthcoming) established sharp convergence rates of CTMC approximation for diffusion models with smooth coefficients but non-smooth payoff functions, and proposed grid design principles to ensure nice convergence behaviors. However, their theoretical analysis fails to obtain sharp convergence rates when model coefficients lack smoothness. Moreover, it is unclear how to design the grid of CTMC to remedy the inferior convergence behaviors resulting from non-smooth model coefficients. In this paper, we introduce new ways for the theoretical analysis of CTMC approximation for general diffusion models with non-smooth coefficients. We prove that convergence of option price is only first order in general. However, strikingly, if all the discontinuous points of the model coefficients and the payoff function are in the midway between two grid points, second order convergence in the maximum norm is restored and in this case, delta and gamma have second order convergence at almost all grid points except those next to the discontinuous points. Numerical experiments are conducted that confirm the validity of our theoretical results. We also compare the CTMC approximation approach with properly designed grids to a classical numerical PDE scheme for diffusion models with non-smooth coefficients, where the finite difference method is applied separately in each region with smooth coefficients and continuous pasting of the value function is enforced at the discontinuities. We show that our approach is superior to the latter in terms of both the convergence rate and the simplicity of implementation.

A General Method for Valuation of Drawdown Risk under Markovian Models

G. Zhang, The Chinese University of Hong Kong (Shenzhen), China, zhanggongqiu@cuhk.edu.cn
L. Li, The Chinese University of Hong Kong, Hong Kong SAR, lfli@se.cuhk.edu.hk

Drawdown measures the price drop from the maximum price and drawdown derivatives are designed to protect investors from sharp price falls. Valuation of these derivatives requires the analysis of the first passage time of the drawdown process, which is a challenging task. Despite the success of some approaches for specific models, a general and computationally efficient method is still lacking. This paper develops a novel algorithm based on continuous time Markov chain approximation and numerical integration to price drawdown derivatives for general Markovian asset price models. Our method allows models with and without jumps to be treated in a unified way. We prove convergence of our algorithm for general Markovian models and provide sharp estimates of the convergence rates for a general class of jump-diffusion models. We also show how to design the grid of the Markov chain to accelerate convergence. Various numerical experiments document the computational efficiency of our method and its advantages over some popular alternatives. In an empirical study, we investigate the effectiveness of a list of drawdown derivatives for hedging the price risk of bitcoin. Our result offers a potential solution to ease concerns over extreme volatility risk in this market.

Thursday 9:00–10:30

Session 4.2 in Room P3 – Asymptotic Behaviour of Stochastic Systems
Chair: Peter Braunsteins

Fast approximate simulation of long-range finite spin systems

R. McVinish, University of Queensland, Australia, r.mcvinish@uq.edu.au
L. Hodgkinson, University of Queensland, Australia, liam.hodgkinson@uqconnect.edu.au

A significant number of models from ecology, epidemiology and statistical mechanics can be formulated as long-range finite spin systems. Exact simulation of large spin systems can be challenging computationally due to small times between transitions. Inspired by Gillespie's tau leaping algorithm, we propose two time discretisation schemes that result in approximate simulation algorithms for finite spin systems. When the transition rates are of a particular commonly used form, fast summation methods can be incorporated into these algorithms resulting in a substantial reduction in computational cost. As the algorithms produce only an approximate simulation, it is important to study the behaviour of the error. We give explicit bounds on the expected number of sites with incorrect spin in terms of the size of the system, time step, and transition rates. In a similar spirit to the error analysis conducted by Anderson, Ganguly, and Kurtz (2011) for the tau leaping algorithm, we also study the behaviour of the rescaled difference between the exact and approximate sample paths.

On the asymptotic behaviour of the number of renewals via translated Poisson

A. Xia, University of Melbourne, Australia, aihuaxia@unimelb.edu.au
Q. Liu, University of Melbourne, Australia, qingweil@student.unimelb.edu.au

In this note, we establish various error estimates for translated Poisson approximation to the number of renewals in terms of Kolmogorov distance, Wasserstein distance and total variation distance. The note complements the error estimate for normal approximation in Kolmogorov distance established in Englund (1980) and improves the bound for compound Poisson approximation in total variation given in Erhardsson (2004).

Performance of the smallest-variance-first rule in appointment sequencing

M. A. de Kemp, University of Amsterdam, the Netherlands, m.a.dekemp@uva.nl
M. Mandjes, University of Amsterdam, the Netherlands, m.r.h.mandjes@uva.nl
N. Olver, Vrije Universiteit Amsterdam, the Netherlands, n.olver@vu.nl

A classical problem in appointment scheduling, with applications in health care, concerns the determination of the patients' arrival times that minimize a cost function that is a weighted sum of mean waiting times and mean idle times. Part of this problem is the sequencing problem, which focuses on ordering the patients. We assess the performance of the smallest-variance-first (SVF) rule, which sequences patients in order of increasing variance of their service durations. While it was known that SVF is not always optimal, many papers have found that it performs well in practice and simulation. We give theoretical justification for these observations by proving quantitative worst-case bounds on the ratio between the cost incurred by the SVF rule and the minimum attainable cost, in a number of settings. We also show that under quite general conditions, this ratio approaches 1 as the number of patients grows large, showing that the svf rule is asymptotically optimal. While this viewpoint in terms of approximation ratio is a standard approach in many algorithmic settings, our results appear to be the first of this type in the appointment scheduling literature.

Local limit theorems for occupancy models

A. Barbour, Universität Zürich, Switzerland, a.d.barbour@math.uzh.ch
P. Braunsteins, The University of Queensland, Australia, p.braunsteins@uq.edu.au
N. Ross, The University of Melbourne, Australia, nathan.ross@unimelb.edu.au

Local central limit theorems for general sums of independent integer valued random variables are well understood; however, for sums of dependent random variables much less is known. In this talk we establish a local central limit theorem for two sums of dependent random variables: the number of degree d vertices in an Erdős–Rényi graph, and the number of germs with d neighbours in a germ-grain model. Our approach relies on Stein's method for distributional approximation.

Session 4.3 in Room P4 – Many-Server Queues
Chair: David Stanford

Monotonicities in systems of parallel Processor Sharing queues

N. Hermansson, University of Auckland, New Zealand, niffe.hermansson@auckland.ac.nz
I. Ziedins, University of Auckland, New Zealand, i.ziedins@auckland.ac.nz

Systems of parallel Processor Sharing queues are useful for modelling various systems ranging from road networks to cloud computing centres. However, it has previously been seen that certain methods for finding user optimal policies for systems of parallel processor sharing queues give results with counterintuitive, and in some ways, problematic non-monotonicities. In order to understand, control, and describe these systems we have analyzed the system behaviour by the method of stochastic coupling and shown a number of important monotonicities to hold in this system, under certain assumptions.
In this talk we give a description of the issues previously found with policy iteration in these systems and then discuss the monotonicity results we have shown to hold. In discussing these results we take special notice of the fact that some are very strong and hold for quite general versions of the system, even in the context of an arbitrary number of queues in parallel, while results that seem quite similar require very stringent constraints on the system in order to hold. We also discuss one result that we would have expected to be able to show using stochastic coupling, but couldn't, and how this failure may relate to the non-monotonicities previously observed in policy iteration in this system.

A symmetric Erlang loss queue with breakdowns

L. van Kreveld, University of Amsterdam, Netherlands, l.r.vankreveld@uva.nl
O. J. Boxma, Einhoven University of Technology, Netherlands, o.j.boxma@tue.nl
J. L. Dorsman, University of Amsterdam, Netherlands, j.l.dorsman@uva.nl
M. Mandjes, University of Amsterdam, Netherlands, m.r.h.mandjes@uva.nl

In this presentation we address an extension of the Erlang loss queue, where the machines (servers) are subject to breakdown. We assume that breakdowns only occur during the service of a product (customer), and that a product cannot be served while its machine is broken. Broken machines are sent to a single-server repair queue, and continue serving its product once it is repaired. This system exhibits interesting relations between the number of busy and broken machines, in such a way that the model can be viewed as a symmetric closed queueing network. We analyze the performance of this network from various perspectives, including different scaling regimes (such as the Halfin-Whitt regime). This analysis provides possibilities for optimization questions.

The Battle between Infinite Server Queues and Heavy-tailed Arrivals

H. Toyoizumi, Waseda University, Japan, toyoizumi@waseda.jp

We find a new stability condition for $M^X/G/\infty$ queue with the heavy-tailed batch size X. Infinite server queues have a great power to process a huge service demand, even when the arrival is so heavy-tailed that its expected demand is infinite. We obtain the condition where the power of infinite server exceeds the heavy-tailed arrival. In addition, by defining simple power-law type distributions, an example of heavy-tailed arrivals, we obtain the probability generating functions of the number of customers when the queue is the stable.

The average waiting time for both classes in a delayed accumulating priority queue

D. A. Stanford, University of Western Ontario, Canada, stanford@stats.uwo.ca
B. Bilodeau, University of Toronto, Canada, blair.bilodeau@mail.utoronto.ca

Mojalal et al (2019) presents the waiting time distribution for the lowest class in the Delayed Accumulating Priority Queue (APQ), in which lower priority class patients wait for a fixed period of time before starting to accumulate priority. Certain simplifying assumptions which apply in the analysis of this case do not extend to the higher-priority classes, which is why analytical expressions for the waiting time distributions for these classes have not yet been found. However, in order to identify the optimal period of delay d before a low priority patient can accumulate priority credit, some quantification of benefit gained by the higher-priority class patients is needed. In the absence of the high-priority waiting time distribution, we focus instead on the exact average waiting times before service for both classes of patients a 2-class Delayed APQ, by making use of a conservation law for average waiting times. We exploit a theorem in Mojalal et al (2019) to quantify the impact upon the average waiting times seen by both classes of patients as a function of the value of the initial delay period d. Numerical investigations in the context of selected KPIs will be presented.

Session 4.4 in Room P5 – Recent Advances in Stochastic Systems I
Chair: Jing Dong

Uniformly Bounded Regret in the Multi-Secretary Problem

A. Arlotto, Duke University, United States, alessandro.arlotto@duke.edu
I. Gurvich, Cornell University, United States, gurvich@cornell.edu

In the secretary problem of Cayley (1875) and Moser (1956), n non-negative, independent, random variables with common distribution are sequentially presented to a decision maker who decides when to stop and collect the most recent realization. The goal is to maximize the expected value of the collected element. In the k-choice variant, the decision maker is allowed to make $k \leq n$ selections to maximize the expected total value of the selected elements. Assuming that the values are drawn from a known distribution with finite support, we prove that the best regret—the expected gap between the optimal online policy and its offline counterpart in which all n values are made visible at time 0—is uniformly bounded in the number of candidates n and the budget k. Our proof is constructive: we develop an adaptive Budget-Ratio policy that achieves this performance. The policy selects or skips values depending on where the ratio of the residual budget to the remaining time stands relative to multiple thresholds that correspond to middle points of the distribution. We also prove that being adaptive is crucial: in general, the minimal regret among non-adaptive policies grows like the square root of n. The difference is the value of adaptiveness.

High-order accuracy steady-state diffusion approximations

J. Dai, The Chinese University of Hong Kong, Shenzhen and Cornell University, jimdai@cuhk.edu.cn

Through numerical work and error bounds via Stein's method on a selected queueing systems, I will demonstrate the benefit of developing diffusion approximations whose diffusion coefficients are state-dependent; these coefficients have been set to be constants in most of the literature in the last 50 years.

Heavy-Traffic Analysis of Queueing Systems with no Complete Resource Pooling

D. Hurtado-Lange, Georgia Institute of Technology, USA, d.hurtado@gatech.edu
S. Mou, Georgia Institute of Technology, USA, smou7@gatech.edu
S. T. Maguluri, Georgia Institute of Technology, USA, siva.theja@gatech.edu

Resource Allocation problems in Stochastic Processing Networks arise in a variety of systems including wired and wireless networks, cloud computing, data centers, manufacturing systems, transportation, call centers, health care etc. Exactly characterizing the performance of these systems is intractable and so, these systems are studied in various asymptotic regimes, heavy-traffic being one of the popular ones. When the so-called complete resource pooling condition is satisfied, these systems usually have a single bottleneck resource in the heavy-traffic limit and so behave like a single server queue. There are many systems such as the Input-queued switch and the bandwidth sharing system that do not satisfy the complete resource pooling, and consequently their heavy-traffic behavior is not completely understood. I will present recent results obtained in such systems using the Drift method. Certain results in non-heavy-traffic regimes will also be presented.

Sub-diffusive load-balancing in time-varying queueing systems

R. Atar, Technion, Israel, rami@ee.technion.ac.il
I. Keslassy, Technion, Israel, isaac@ee.technion.ac.il
G. Mendelson, Technion, Israel, galmen@campus.technion.ac.il

The degree to which delays or queue lengths equalize under load balancing algorithms gives a good indication on their performance. Some of the most well-known results in this context are concerned with the asymptotic behavior of the delay or queue length at the diffusion scale under a critical load condition, where arrival and service rates do not vary with time. For example, under the join the shortest queue policy, the queue length deviation process, defined as the difference between the greatest and smallest queue length as it varies over time, is at a smaller scale (sub-diffusive) than that of queue lengths (that is diffusive). Denoting by n the usual scaling parameter, the former and the latter are of order $o(n^{1/2})$ and $O(n^{1/2})$, respectively.
The goal of this talk is to argue that sub-diffusivity of the deviation process (SDDP) is not limited to heavy traffic. To exhibit the breadth of this phenomenon, we establish SDDP in three load-balancing models: power of d choices, redundancy routing and longest queue first. Our results accommodate server heterogeneity and time-varying arrival and service intensities, which cause transitions between underloaded, critically loaded and overloaded regimes. As an application of these results, asymptotic optimality of the first two policies is shown, with an optimality guarantee of order $o(n^{1/2})$ in the aforementioned framework, where in particular queue sizes may reach $O(n)$.

Session 4.5 in Room M1 – New Developments in Reinforcement Learning
Chair: Mengzhou Liu

Thompson Sampling with Information Relaxation Penalties

S. Min, Columbia Business School, U.S., smin20@gsb.columbia.edu
C. Maglaras, Columbia Business School, U.S., c.maglaras@gsb.columbia.edu
C. C. Moallemi, Columbia Business School, U.S., ciamac@gsb.columbia.edu

We consider a finite time horizon multi-armed bandit (MAB) problem in a Bayesian framework, for which we develop a general set of control policies that leverage ideas from information relaxations of stochastic dynamic optimization problems. In crude terms, an information relaxation allows the decision maker (DM) to have access to the future (unknown) rewards and incorporate them in her optimization problem to pick an action at time t, but penalizes the decision maker for using this information. In our setting, the future rewards allow the DM to better estimate the unknown mean reward parameters of the multiple arms, and optimize her sequence of actions. By picking different information penalties, the DM can construct a family of policies of increasing complexity that, for example, include Thompson Sampling and the true optimal (but intractable) policy as special cases.
We systematically develop this framework of *information relaxation sampling*, propose an intuitive family of control policies for our motivating finite time horizon Bayesian MAB problem, and prove associated structural results and performance bounds. Numerical experiments suggest that this new class of policies performs well, in particular in settings where the finite time horizon introduces significant tension in the problem. Finally, inspired by the finite time horizon Gittins index, we propose an index policy that builds on our framework that particularly outperforms to the state-of-the-art algorithms in our numerical experiments.

Continuous-Time Mean-Variance Portfolio Selection: A Reinforcement Learning Framework

H. Wang, Columbia University, USA, hw2718@columbia.edu
X. Zhou, Columbia University, USA, xz2574@columbia.edu

We approach the continuous-time mean–variance (MV) portfolio selection with reinforcement learning (RL). The problem is to achieve the best tradeoff between exploration and exploitation, and is formulated as an entropy-regularized, relaxed stochastic control problem. We prove that the optimal feedback policy for this problem must be Gaussian, with time-decaying variance. We then establish connections between the entropy-regularized MV and the classical MV, including the solvability equivalence and the convergence as exploration weighting parameter decays to zero. Finally, we prove a policy improvement theorem, based on which we devise an implementable RL algorithm. We find that our algorithm outperforms both an adaptive control based method and a deep neural networks based algorithm by a large margin in our simulations.

Duality-Based Exploration in Reinforcement Learning and Its Application

N. Chen, The Chinese University of Hong Kong, Hong Kong, nchen@se.cuhk.edu.hk
X. Ma, The Chinese University of Hong Kong, Hong Kong, xma@se.cuhk.edu.hk

We develop a Boltzmann exploration approach to address the exploration-and-exploitation tradeoff in reinforcement learning problems. Based on a new concept of information duality, the proposed approach will probabilistically choose under-explored states to guide exploration. Numerical implementations in some financial applications demonstrates the efficiency of the method.

Dynamic Pricing of Relocating Resources in Large Networks

S. R. Balseiro, Graduate School of Business, Columbia University, United States, srb2155@columbia.edu
D. B. Brown, Fuqua School of Business, Duke University, United States, dbbrown@duke.edu
C. Chen, Fuqua School of Business, Duke University, United States, cc459@duke.edu

Motivated by applications in shared vehicle systems, we study dynamic pricing of resources that relocate over a network of locations. Customers with private willingness-to-pay sequentially request to relocate a resource from one location to another, and a revenue-maximizing service provider sets a price for each request. This problem can be formulated as an infinite horizon stochastic dynamic program, but is quite difficult to solve, as optimal pricing policies may depend on the locations of all resources in the network. We first focus on networks with a hub-and-spoke structure, and we develop a dynamic pricing policy and a performance bound based on a Lagrangian relaxation. This relaxation decomposes the problem over spokes and is thus far easier to solve than the original problem. We analyze the performance of the Lagrangian-based policy and focus on a *supply-constrained large network* regime in which the number of spokes (n) and the number of resources grow at the same rate. We show that the Lagrangian policy loses no more than $O\big(\sqrt{\ln n/n}\big)$ in performance compared to an optimal policy, thus implying asymptotic optimality as n grows large. We also show that the Lagrangian upper bounds are tighter than upper bounds from fluid relaxations, and provide examples that show that the fluid-based upper bounds and their induced static policies fail to be asymptotically optimal. Finally, we extend the approach to general networks with multiple, interconnected hubs and spoke-to-spoke connections. For the special case with uniformly related hubs - networks in which the ratio of the arrival rate to a spoke from a hub to the arrival rate of the reverse trip is constant across hubs - we provide analogous performance bounds and show that the policy loses no more than $O\big(\max\big\{\sqrt{J/n}, \sqrt{\ln n/n}\big\}\big)$ in performance in the large network regime when the number of hubs J grows at any rate $o(n)$.

Session 4.6 in Room M2 – Diffusions and Related Stochastic Processes
Chair: Jiyeon Lee

Probabilistic Contraction Analysis of Iterated Random Operators

A. Gupta, Ohio State University, gupta.706@osu.edu
R. Jain, University of Southern California, rahul.jain@usc.edu
P. Glynn, Stanford University, glynn@stanford.edu

Any stochastic iterative algorithm may be seen as iteration of a random function or operator. Indeed, this motivated Diaconis and Freedman (1999) to study iterated random functions. Iteration of deterministic operators is usually studied by either using Lyapunov techniques or using their contraction properties. Iteration of random operators when they occur in stochastic iterative algorithms is either studied via Stochastic Lypunov theory, or by using direct properties of the related stochastic processes, e.g., Markov or Martingale property. Indeed, Diaconis and Freedman used the Markov property to argue convergence and find rate of convergence. But in many problems, this can be quite cumbersome, while constructing Stochastic Lyapunov functions is much harder than constructing Lyapunov functions for deterministic settings. In this paper, we develop rudiments of the technique of probabilistic contraction analysis of iterated random operators. This is applicable when the random operator has some sort of a probabilistic contraction property. We show a number of ways this property can be defined. Accordingly, we define notions of probabilistic fixed points of such random operator in the spirit of the Banach fixed point theorem for deterministic contracting operators. We then prove that iteration of such random operators indeed converges to such probabilistic fixed points. Such random operators may arise in a variety of situations. For example, in dynamic programming algorithms, the Bellman operator may be replaced by an empirical Bellman operator obtained by doing sample-average approximation of the expectation in the operator. This works quite well numerically in both the discounted setting as shown in Haskell, Jain, and Kalathil (2016) as well as the average setting. We can argue that any stochastic approximation scheme fits within this framework. In fact, we show that even algorithms for stochastic optimization such as Batch Stochastic Gradient Descent (B-SGD) as well as SVRG (SGD with Variance Reduction) may also be viewed in this way.
We argue that a key appeal of the probabilistic contraction analysis method is that it is quite easy to operationalize by constructing a Markov Chain that stochastically dominates the error process. This in most problems is quite natural. Then, the convergence analysis just reduces to analysis of a finite state Markov Chain which we can show has an invariant distribution that concentrates at zero. In the future, we will build on this framework to increase the applicability of this analysis technique to a larger class of stochastic iterative algorithms. The goal is a technique that is easy to operationalize, allows us to conclude convergence with reasonable effort, and also yields non-asymptotic rate of convergence.

Reflecting diffusions in nonsmooth domains: Markov selection and uniqueness

C. Costantini, Università di Chieti-Pescara, Italy, c.costantini@unich.it
T. G. Kurtz, University of Wisconsin-Madison, USA, kurtz@math.wisc.edu

The issue of existence and uniqueness of obliquely reflecting diffusions is still open for many nonsmooth domains that arise, for instance, in queueing networks approximation.
A successful approach to proving uniqueness (adopted, for instance, in Dai and Williams (1996) and Costantini and Kurtz (2018)) relies on existence of a Markov selection, i.e. a Markov process satisfying the definition of reflecting diffusion. Existence of such a Markov selection is not an immediate generalization of the corresponding results for unconstrained diffusions and requires new tools.
We prove existence of a Markov selection for obliquely reflecting diffusions under very general conditions on the domain and the direction of reflection.
Our result allows to prove some new uniqueness results. For instance, we prove uniqueness in a 2-dimensional piecewise smooth domain, for directions of reflection that do not satisfy the conditions of Dupuis and Ishii (1993).

The tilted flashing Brownian ratchet

S. N. Ethier, University of Utah, USA, ethier@math.utah.edu
J. Lee, Yeungnam University, South Korea, leejy@yu.ac.kr

The flashing Brownian ratchet is a stochastic process that alternates between two regimes, a one-dimensional Brownian motion and a Brownian ratchet, the latter being a one-dimensional diffusion process that drifts towards a minimum of a periodic asymmetric sawtooth potential. The result is directed motion. In the presence of a static homogeneous force that acts in the direction opposite that of the directed motion, there is a reduction (or even a reversal) of the directed motion effect. Such a process may be called a tilted flashing Brownian ratchet. We show how one can study this process numerically, using a random walk approximation.

How strong can the Parrondo effect be?

S. N. Ethier, University of Utah, USA, ethier@math.utah.edu
J. Lee, Yeungnam University, South Korea, leejy@yu.ac.kr

The Parrondo effect appears when two fair coin-tossing games played in a random sequence or in some periodic sequence form a winning game. In the capital-dependent games of Parrondo, game A uses a fair coin, while game B uses two biased coins, a p_0-coin if current capital is divisible by r and a p_1-coin otherwise. In both games, the player wins one unit with heads and loses one unit with tails. Game B is fair if $(1-p_0)(1-p_1)^{r-1} = p_0 p_1^{r-1}$. We prove a strong law of large numbers for arbitrary periodic sequences of Parrondo games. Using these results we show that if the parameters of game B are allowed to be arbitrary, subject to the fairness constraint, and if the two (fair) games A and B are played in an arbitrary periodic sequence, then the rate of profit can not only be positive, it can be arbitrarily close to 1 (i.e., 100%).

Session 4.7 in Rooms M5 and M6 – Optimal Stopping and Control
Chair: David Goldberg

A Finite Time Analysis of Temporal Difference Learning

J. Bhandari, Columbia Univeriesty, jb3618@columbia.edu
D. Russo, Columbia, University, djr2174@gsb.columbia.edu
R. Singal, Columbia University, rs3566@columbia.edu

Temporal difference learning (TD) is a simple iterative algorithm used to estimate the value function corresponding to a given policy in a Markov decision process. Although TD is one of the most widely used algorithms in reinforcement learning, its theoretical analysis has proved challenging and few guarantees on its statistical efficiency are available. In this work, we provide a simple and explicit finite time analysis of temporal difference learning with linear function approximation. Except for a few key insights, our analysis mirrors standard techniques for analyzing stochastic gradient descent algorithms, and therefore inherits the simplicity and elegance of that literature. Final sections of the paper show how all of our main results extend to the study of TD learning with eligibility traces, known as TD(λ), and to Q-learning applied in high-dimensional optimal stopping problems.

A new approach to high-dimensional online decision-making

Y. Chen, Cornell University, USA, yc2436@cornell.edu
D. A. Goldberg, Cornell University, USA, dag369@cornell.edu

Big data promotes the use of more realistic models to describe sequential decision-making in many practical settings, such as options pricing, dynamic pricing and revenue management, personalized medicine, and operations more broadly. These models track the dynamics of complex features (e.g. historical data, economic indicators, or personalized information) that can affect the optimal decision over time, often boiling down to solving high-dimensional stochastic dynamic programs (DP). However, such high-dimensional stochastic DPs are well known to be computationally intractable in general, suffering from the so-called “curse of dimensionality”. To sidestep this fundamental computational issue, much of the literature either makes restrictive assumptions on the underlying model and/or distributions/ transitions of the DP, or resorts to approximation methods that are mostly heuristic in nature.

We build on recent algorithmic progress of Chen and Goldberg for optimal stopping to devise a new framework that overcomes this fundamental computational barrier for a fairly general class of such high-dimensional stochastic DP, including the aforementioned and many other cases. Our results come with strong theoretical guarantees (both runtime and accuracy), subject to a condition we refer to as the “limited policy changes” (LPC) property, which (for example) in the dynamic pricing setting translates to a “limited price experimentation" assumption; and in the multiple stopping setting translates to a “limited number of stops" assumption. Beyond this natural constraint, our results allow for arbitrarily complex and non-stationary high-dimensional dependencies and features, assuming only access to a black-box simulator for the underlying stochastic processes. Time permitting, we also discuss several improvements to the original approach of Chen and Goldberg for optimal stopping.

Beating the curse of dimensionality in options pricing and optimal stopping

D. A. Goldberg, Cornell University, USA, dag369@cornell.edu
Y. Chen, Cornell University, USA, yc2436@cornell.edu

The fundamental problems of pricing high-dimensional path-dependent options and optimal stopping are central to applied probability, financial engineering, operations research, and stochastic control. Modern approaches, often relying on ADP, simulation, and/or duality, typically have limited rigorous guarantees, which may scale poorly and/or require previous knowledge of good basis functions. A key difficulty with many approaches is that to yield stronger guarantees, they would necessitate the computation of deeply nested conditional expectations, with the depth scaling with the time horizon T.
We overcome this fundamental obstacle by providing an algorithm which can trade-off between the guaranteed quality of approximation and the level of nesting required in a principled manner. We develop a novel pure-dual approach, inspired by a connection to network flows. This leads to a representation for the optimal value as an infinite sum for which: 1. each term is the expectation of an elegant recursively defined infimum; 2. the first k terms only require k levels of nesting; and 3. truncating at the first k terms yields a (normalized) error of $1/k$. This enables us to devise simple randomized and data-driven algorithms and stopping strategies whose runtimes are effectively independent of the dimension, beyond the need to simulate sample paths of the underlying process. Our method allows one to elegantly trade-off between accuracy and runtime through a parameter epsilon controlling the associated performance guarantee (analogous to the notion of PTAS in the theory of approximation algorithms), with computational and sample complexity both polynomial in T (and effectively independent of the dimension) for any fixed epsilon, in contrast to past methods typically requiring a complexity scaling exponentially.

McKean stochastic optimal control of an energy storage system to reduce demand variability

M. Grangereau, EDF and CMAP, Ecole polytechnique, France, maxime.grangereau@edf.fr
E. Gobet, CMAP, Ecole polytechnique, France, emmanuel.gobet@polytechnique.edu

We consider a system, e.g., a micro-grid, characterized by its own exogenous production and consumption, connected to the network and to a controllable energy storage system, like a battery for instance. We control this energy storage system and aim at limiting simultaneously the power demand peaks on the network, the storage system aging and the fluctuations of the power supplied by the network. While the first two components of the cost are related to the trajectory of the system, and therefore to the standard stochastic optimal control paradigm, the last one is more related to the probability distribution of the system, and therefore to the McKean stochastic optimal control paradigm.
We propose here to model the problem of energy storage control as a McKean stochastic optimal control problem with scalar interactions. A certain number of theoretical results can be obtained for this problem: necessary and sufficient optimality conditions, uniqueness and existence of a solution. We focus here on two particular cases of the problem where a solution can be computed very efficiently. For the so-called Linear-Quadratic case, some quasi-explicit formulas are available for the solution of the control problem, and solving the problem amounts to solving two Riccati Ordinary Differential Equations and two Backward Stochastic Differential Equations. In the non-quadratic case, a first order expansion of the optimal control with respect to the non-quadratic terms can be obtained based on a perturbation approach of the linear-quadratic case. The performance of our approach is demonstrated through numerical examples in both the linear-quadratic and quasi linear-quadratic cases.

Session 4.8 in Rooms M7 and M8 – Healthcare I
Chair: Zeynep Akşin Karaesmen

Balancing Multi-Featured Workload in Service Systems

N. Carmeli, Technion—Israel Institute of Technology, Israel, nitzany@campus.technion.ac.il
G. B. Yom-Tov, Technion—Israel Institute of Technology, Israel, gality@technion.ac.il
A. Mandelbaum, Technion—Israel Institute of Technology, Israel, avim@technion.ac.il

Operational offered-load of servers, and its calculated agents utilization through queueing theory, support the design, staffing, and routing of congestion-prone service systems. However, in many such systems, and healthcare in particular, service providers experience additional features of load such as emotional and cognitive; these type of load may cause stress and fatigue even when the offered load is low. How does one balance such load among servers, in a way that accommodates its multi-features? In our research we develop a theoretical framework that supports answers to this question. More concretely, we consider a queueing system with several agent skills/groups and customer classes (e.g. skills-based routing in call centers), within which we optimally and fairly balance multi-feature load.

Sizing flexible resources in large-scale service systems

J. Chen, Columbia University, USA, jc4823@columbia.edu
J. Dong, Columbia University, USA, jing.dong@gsb.columbia.edu

Service systems typically involve multiple classes of customers, each requiring a different type of service. Furthermore, the servers are often classified into different pools, each with different skill requirements. For example, in call centers, customers may require types of services, and agents are trained in one or several of these specialized services. From the system design point of view, increased server flexibility can improve system performance by allowing an underloaded server pool to help an overloaded pool during imbalances that arise due to stochastic fluctuations. This is known as resource pooling. However, flexible servers are typically more costly than dedicated servers (servers who can only serve one specific customer class), and moreover may provide service at a slower rate. It is thus unclear what the optimal degree of flexibility is. Our work aims to determine the right sizing for flexible resources in large-scale service systems. We find that the optimal size of flexible resources is always much smaller (order of magnitude smaller) than that of dedicated resources, and that the exact size depends on the cost difference and the service quality. When cost and service quality differences are asymptotically negligible, we will staff the flexible resource in a regime where we achieve complete resource pooling. Otherwise, we will staff the flexible resource in a regime where we achieve only partial resource pooling.

Behavior of strategic patients in a health care system with public and private facilities

L. Örmeci, Koç University, Turkey, lormeci@ku.edu.tr
D. Andritsos, HEC, France, andritsos@hec.fr
P. Canbolat, Koç University, Turkey, pcanbolat@ku.edu.tr
Y. Dimitrakopoulos, University of Athens, Greece, dimgiannhs81@gmail.com
S. Sadeghzadeh, Virginia Tech University, USA, saloumeh@vt.edu

In most countries, health care facilities operate in both private and public sectors, where patients choose to receive service from either a public or private facility to maximize his/her expected utility. Hence, the behavior of patients determines the average waiting times of patients and the utilization of facilities in the two sectors, which, in turn, affects the associated service fees and capacity levels. In this talk, we consider two settings: the first one focuses on diagnostic services, whereas the other one models a dual-practice physician. We characterize the patient behavior using queueing games and analyze its effects on the system operations in both settings.

Modelling intensive care units using quasi-birth-death processes

S. James, University of Adelaide, Australia, sarah.james@adelaide.edu.au
N. Bean, University of Adelaide, Australia, nigel.bean@adelaide.edu.au
J. Tuke, University of Adelaide, Australia, simon.tuke@adelaide.edu.au

An intensive care unit (ICU) is a crucial and limited resource in a hospital which is affected by uncertainty and variability, and often operates close to capacity. Queueing theory has been used to model the bed occupancy in ICUs for the last 20 years, with a particular focus on using $M/M/\cdot$ and $M/PH/\cdot$ queueing models. These queueing models assume that the arrival process (patient arrivals) is a Poisson process, the service times (length of stay) follow an exponential or Phase-Type distribution, and most crucially that the arrival process and service times are independent of each other. However, Varney *et al.* (2018) showed that there is some dependence structure between the arrival process and the service times in an ICU. Without independence between the arrival process and service times, standard queueing models become invalid.
We aim to provide a more principled approach to modelling bed occupancy in ICUs using quasi-birth-and-death processes (QBDs). By allowing the phases of the arrival process and the service times to interact with each other, QBDs provide the flexibility to model a queueing system with dependence between the arrival process and the service times. In this talk, we describe the approaches we have taken to develop model-fitting procedures for QBD models, with a particular focus on fitting suitable QBD models to data from an intensive care unit.

Thursday 13:00–14:30

Session 5.1 in Rooms P1 and P2 – Simulation, Risk, and Optimization
Chair: Raghu Pasupathy

An Envelope Procedure for Ranking and Selection Problems

S. Ma, Cornell University, Ithaca
S. G. Henderson, Cornell University, Ithaca, sgh9@cornell.edu

Ranking and Selection (R & S) procedures are commonly used for choosing the best among a finite set of candidate systems, each of which is associated with a simulation model. We present a new fully sequential selection procedure, the Envelope Procedure (EP), that provides a probably-approximately-correct guarantee on the selection. The EP has a rigorous complexity analysis for a particular choice of sampling strategy. Numerical experiments show that EP is more efficient than some leading R & S procedures in many problem configurations.

Systemic Risk - A simulation perspective

A. Deo, TIFR, Mumbai, India, anand.deo@tifr.res.in
S. Juneja, TIFR, Mumbai, India, juneja@tifr.res.in

Post the crisis of 2008, systemic risk, or the cascading failures among banks in an economy due to distressed assets of a few banks has become a central area of research. Banks an economy borrow and lend capital among each other, and it is of vital importance to identify systemically important banks, as well as understand the global properties of the network itself (for example average profit made by banks or fraction of banks which default on their payments). This is complicated by the fact that the modeller typically only observes partial information about the network.
In this talk, we take a simulation approach to the above problem. Modelling the banking network as a marked random graph whose statistical properties match those of the observed network, we show that as the size of the economy increases, the wealths of the banks in presence of partial information converge to the fixed point of a distribution valued map. Using this, we infer that the global properties of the network are robust to the actual realisation of partial information, and only depend on the limiting fixed point. The modeller may then approximate a large fraction of banking networks by simulating this limiting fixed point distribution. To this end, we develop a Monte-Carlo algorithm for computing fixed points of distribution valued maps, and provide rigorous guarantees on its performance. This allows us to compute the total profit, as well as identify systemically important banks: banks which if they fail cause large losses to the economy.

Convergence Rates of Sampling Algorithms for Stochastic Optimization: The Effect of Curvature

D. Newton, Purdue University, Department of Statistics, newton34@purdue.edu
R. Pasupathy, Purdue University, Department of Statistics, pasupath@purdue.edu
A. N. K. Yip, Purdue University, Department of Mathematics, yipn@purdue.edu

The Stochastic Approximation (SA) recursion $X_{k+1} = X_k - \gamma_k \nabla F(X_k), k = 1, 2, \ldots$, also known as Stochastic Gradient Descent (SGD), is the "workhorse" stochastic optimization recursion for machine learning. For smooth, strongly convex objectives, it is well-known that the SA iteration achieves the information theoretic Cramér-Rao lower bound (not to be confused with the more common but crude optimal rate $\mathcal{O}(k^{-1/2})$) when the step size $\gamma_k := \theta/k$ and $\theta = c^{-1}$, where c is the smallest eigen value of the Hessian of the objective function at the optimum. When $\theta < (2c)^{-1}$, SA's convergence rate deteriorates rapidly, failing to achieve even the crude $\mathcal{O}(k^{-1/2})$ rate. There is strong indication that this situation is not merely "theoretical" but manifests frequently in practice because the eigen values of the Hessian at the optimum are rarely known in advance. Moreover, practitioners vastly prefer using a fixed-step recursion owing to fast transience. We ask whether these riddles can be remedied through sampling. We thus construct a fixed-step sampling recursion where the inverse Hessian and gradient estimates are constructed using sampling, and potentially updated on different timescales. A fine analysis of the resulting recursion reveals that the Cramér-Rao information-theoretic bound is attained to arbitrary accuracy with only a "light updating" of the Hessian inverse, as long as the gradient estimation sample sizes increase at a specific rate. Unlike in optimal SGD, none of the optimal parameter prescriptions within the proposed procedure depend on unknown curvature constants. To fully illuminate the effect of (a) dimension, (b) computation, and (c) sampling, we express all convergence rates in terms of a work measure that includes all such costs. Our numerical results are consistent with success reported by a number of recent adaptive sampling heuristics.

Optimal Transport Based Distributionally Robust Optimization: Structural Properties and Iterative Schemes

J. Blanchet, Stanford University, USA, jose.blanchet@stanford.edu
K. Murthy, SUTD, Singapore, karthyek_murthy@sutd.edu.sg
F. Zhang, Stanford University, USA, fzh@stanford.edu

We consider optimal transport based distributionally robust optimization (DRO) problems with locally strongly convex transport cost functions and affine decision rules. We obtain structural results about the value function, the optimal policy, and the worst-case optimal transport adversarial model. Our results expose a rich structure which then we use to obtain optimization procedures which have basically the same sample and iteration complexity as non-DRO stochastic gradient descent and sometimes even better complexity as such a benchmark.

Session 5.2 in Room P3 – Stochastic Network Asymptotics
Chair: Rami Atar

JSQ, LDP, and the Golden Ratio

A. Budhiraja, University of North Carolina at Chapel Hill, NC 27599-3260, budhiraj@email.unc.edu
E. Friedlander, University of Chicago, Chicago, IL 60637, efriedlander@uchicago.edu
R. Wu, University of Michigan Ann Arbor, Ann Arbor, MI, 48109, ruoyu@umich.edu

Consider n urns such that each of them contains exactly one ball at time 0. Balls arrive at rate n with exponential interarrival times and an arriving ball is dropped into an urn with the fewest balls (with ties broken at random). Balls depart from a nonempty urn at rate 1 according to an exponential distribution. All interarrival times and departure times are mutually independent. This is known as the Join-the-Shortest-Queue (JSQ) system. Let $X_i^n(t)$ denote the proportion of urns with i or more balls at time t. We establish a large deviation principle for the occupancy process $X^n = (X_1^n, X_2^n, \cdots)$ in a suitable path space. Using this LDP we then give exponential decay rates for probabilities of long queues. The Golden Ratio makes an appearance.

Heavy traffic limits for load balancing algorithms using delayed information

R. Atar, Technion — Israel Institute of Technology, Israel, rami@technion.ac.il
D. Lipshutz, Technion — Israel Institute of Technology, Israel, lipshutz@technion.ac.il

We consider a load balancing problem for a network of parallel queues in which incoming jobs (or customers) are routed to one of the queues upon arrival. In our model, the routing decision depends only on past routing decisions and the delayed state of the queues. When a job arrives to the queue, the current state of the queues is estimated based on this information and the incoming job is routed to the queue with the shortest estimated length. We formulate a setting where the delay remains constant under heavy traffic scaling, so the effects of the delay do not disappear. We prove state space collapse of the estimated queue lengths under heavy traffic scaling. This allows us to formulate a diffusion model and prove convergence to it under several natural estimation schemes.

Interacting Markov chains on Large Sparse Graphs

A. Ganguly, Brown University, USA, Ankan_Ganguly@brown.edu
D. Lacker, Columbia University, USA, Daniel.Lacker@columbia.edu
K. Ramanan, Brown University, USA, Kavita_Ramanan@brown.edu
R. Wu, U. Michigan, USA, ruoyu@umich.edu

We consider a large system of homogeneous interacting Markov chains on a (possibly random) graph in which the infinitesimal evolution of each particle depends on its own state and the empirical measure of the states of neighboring particles. Such interacting particle systems arise in a variety of applications, ranging from engineering to physics. When the graph is complete, it is well known that the dynamics of a typical particle converges in the limit, as the number of vertices goes to infinity, to a so-called nonlinear Markov process, often referred to as the McKean-Vlasov or mean-field limit. In this talk, we focus on the complementary case of scaling limits of dynamics on sparse graphs, and obtain a novel characterization of the dynamics of a typical particle in the case when the limit interaction graph is a tree. We also present numerical results to demonstrate the efficacy of using the limit as an approximation for a finite collection of interacting chains.

Logarithmic Regret in the Dynamic and Stochastic Knapsack Problem with Equal Rewards

A. Arlotto, Duke University, USA, alessandro.arlotto@duke.edu
X. Xie, Duke University, USA, xinchang.xie@duke.edu

We study a dynamic and stochastic knapsack problem in which a decision maker is sequentially presented items arriving according to a Bernoulli process over n discrete time periods. Items have equal rewards and independent weights that are drawn from a known continuous distribution F. The decision maker seeks to maximize the expected total reward of items that she includes in the knapsack while satisfying a capacity constraint, and while making terminal decisions as soon as each item weight is revealed. Under mild regularity conditions on the weight distribution F, we prove that the regret—the expected difference between the performance of the best sequential algorithm and that of a prophet who sees all of the weights before making any decision—is, at most, logarithmic in n. Our proof is constructive. We devise a re-optimized heuristic that achieves this regret bound. We also discuss the limiting distribution of the total reward collected by this re-optimized heuristic.

Session 5.3 in Room P4 – Queueing, Control, and Services
Chair: Amy Ward

A Moment Matching Algorithm for Dynamic Programs

I. Gurvich, Cornell University, NY, USA, gurvich@cornell.edu
B.-Z. Zhang, Cornell University, NY, USA, bz275@cornell.edu

We offer an approach to approximate dynamic programming that is based on the creation of an approximate Markov chain whose local transition moments match those of the original chain but whose optimization is computationally more efficient. Specifically, we rely on the well-known soft-aggregation algorithm and optimize its parameters — the aggregation and disaggregation probabilities — to match said local moments. Recent work allows us to “couple” the original chain with the one that we artificially created to derive optimality gap bounds.

Computational Algorithms for Optimal Design of Service Systems

P. Jaiswal, Purdue University, West Lafayette IN, jaiswalp@purdue.edu
H. Honnappa, Purdue University, West Lafayette IN, honnappa@purdue.edu
V. A. Rao, Purdue University, West Lafayette IN, varao@purdue.edu

We consider the problem of optimally designing a service system in a data-driven setting. As an example, consider a staffing problem that can be solved either using an inventory model or a queueing model. Such optimal design problems can be cast as stochastic optimization programs over parametrized stochastic models. In a data-driven setting, where the parameters may not be known fully, a reasonable approach to solving the optimal design problem is to use a *forecast-then-optimize* approach: a predictive model of the system state is inferred from data and incorporated into the stochastic optimization program.
In this talk, we present a general framework for solving such data-driven optimal design problems in a Bayesian setting. Our framework is based on dual representations of, and variational approximations to, the stochastic optimization program. We extract three computational algorithms from this framework, and characterize the *optimality gap* of these methods in terms of information-theoretic finite sample bounds. We illustrate our theory through simulation and real-world example data sets.

Time-Varying Robust-Queueing for Quantile Approximation

W. Whitt, Columbia University, USA, ww2040@columbia.edu
W. You, Columbia University, USA, wy2225@columbia.edu

We develop a time-varying robust-queueing (TVRQ) algorithm for the continuous-time workload in a single-server queue with time-varying arrival-rate and service-rate functions. We examine the performance of TVRQ in predicting the periodic steady-state workload in models with periodic rate functions. We show that the periodic TVRQ converges to a proper limit in appropriate long-cycle and heavy-traffic regimes and coincides with long-cycle fluid limits and heavy-traffic diffusion limits for long cycles. We find that the algorithm predicts the quantiles and the mean of the periodic distribution of the steady-state workload remarkably well. Simulation studies also show that PRQ accurately predicts the time lag of the peak congestion in compare to the peak traffic intensity.

Timing it Right: Balancing Inpatient Congestion versus Readmission Risk at Discharge

P. Shi, Purdue University, USA, shi178@purdue.edu
J. Helm, Indiana University, USA, helmj@indiana.edu
J. Deglise-Hawkinson, Lean Care Solutions Corporation Pte. Ltd
J. Pan, Lean Care Solutions Corporation Pte. Ltd

When to discharge a patient plays an important role in hospital patient flow management and patient outcomes. In this work, we develop and implement a practical decision support tool to aid hospitals in managing the delicate balance between readmission risk at discharge and ward congestion. We formulate the discharge decision framework as a large-scale Markov Decision Process (MDP) that integrates a personalized readmission prediction model to dynamically prescribe both how many and which patients to discharge each day. We overcome challenges from both the analytical and prediction sides. Due to patient heterogeneity and the fact that length-of-stay is not memoryless, the MDP suffers the curse of dimensionality. We derive useful structural properties and leverage an analytical solution for a special cost setting to transform the MDP into a univariate optimization; this leads to an efficient dynamic heuristic. Meanwhile, off-the-shelf prediction models alone could not provide adequate input for our decision support framework. To bridge this gap, we integrate several statistical methods to build a new readmission prediction model that allows us to implement our decision framework with existing hospital data systems. Through extensive counterfactual analyses, we demonstrate the value of our recommended discharge policy over our partner hospital's historical discharge behavior. We also discuss the implementation efforts of this discharge optimization tool at our partner hospital.

Session 5.4 in Room P5 – Recent Advances in Stochastic Systems II
Chair: Alessandro Arlotto

Efficient random graph matching via degree profiles

J. Ding, University of Pennsylvania, USA, dingjian@wharton.upenn.edu
Z. Ma, University of Pennsylvania, USA, zongming@wharton.upenn.edu
Y. Wu, Yale University, USA, yihong.wu@yale.edu
J. Xu, Duke University, USA, jiaming.xu868@duke.edu

Random graph matching refers to recovering the underlying vertex correspondence between two random graphs with correlated edges; a prominent example is when the two random graphs are given by Erdős-Rényi graphs $G(n, \frac{d}{n})$. This can be viewed as an average-case and noisy version of the graph isomorphism problem. Under this model, the maximum likelihood estimator is equivalent to solving the intractable quadratic assignment problem. This work develops an $\tilde{O}(nd^2+n^2)$-time algorithm which perfectly recovers the true vertex correspondence with high probability, provided that the average degree is at least $d = \Omega(\log^2 n)$ and the two graphs differ by at most $\delta = O(\log^{-2}(n))$ fraction of edges. For dense graphs and sparse graphs, this can be improved to $\delta = O(\log^{-2/3}(n))$ and $\delta = O(\log^{-2}(d))$ respectively, both in polynomial time. The methodology is based on appropriately chosen distance statistics of the degree profiles (empirical distribution of the degrees of neighbors). Before this work, the best known result achieves $\delta = O(1)$ and $n^{o(1)} \leq d \leq n^c$ for some constant c with an $n^{O(\log n)}$-time algorithm and $\delta = \tilde{O}((d/n)^4)$ and $d = \tilde{\Omega}(n^{4/5})$ with a polynomial-time algorithm.

The Power of Two in Queue Scheduling

Y. Chen, Columbia University, USA, yc3107@columbia.edu
J. Dong, Columbia University, USA, jing.dong@gsb.columbia.edu

Utilizing customer service time information, we study a scheduling policy with two priority classes. In the single-server heavy-traffic regime, our scheduling policy achieves similar scaling for the queue length processes as the shortest remaining processing time first policy. Our analysis quantifies on how the tail of the service time distribution will affect the benefit one can gain from smart scheduling policies. When the service time information is estimated/predicted, we further analyze how prediction errors will affect the performance our scheduling policy. Our results provide insights on the interplay between the service time distribution and the estimation error distribution on system performance. The imperfect information analysis also demonstrates the robustness of our scheduling policy.

Query Complexity of Bayesian Private Learning

K. Xu, Stanford Graduate School of Business, Stanford, CA, USA, kuangxu@stanford.edu

We study the query complexity of Bayesian Private Learning: a learner wishes to locate a random target within an interval by submitting queries, in the presence of an adversary who observes all of her queries but not the responses. How many queries are necessary and sufficient in order for the learner to accurately estimate the target, while simultaneously concealing the target from the adversary?
Our main result is a query complexity lower bound that is tight up to the first order. We show that if the learner wants to estimate the target within an error of ε, while ensuring that no adversary estimator can achieve a constant additive error with probability greater than $1/L$, then the query complexity is on the order of $L\log(1/\varepsilon)$, as $\varepsilon \to 0$. Our result demonstrates that increased privacy, as captured by L, comes at the expense of a *multiplicative* increase in query complexity.
Our proof method builds on Fano's inequality and a family of proportional-sampling estimators. As an illustration of the method's wider applicability, we generalize the complexity lower bound to settings involving high-dimensional linear query learning and partial adversary observation.

Optimal Contract for Machine Repair and Maintenance

F. Tian, University of Michigan, ftor@umich.edu
P. Sun, Duke University, psun@duke.edu
I. Duenyas, University of Michigan, duenyas@umich.edu

A principal hires an agent to repair a machine when it is down and maintain it when it is up, and earns a flow revenue when the machine is up. Both the up and down times follow exponential distributions. If the agent exerts effort, the downtime is shortened, and uptime is prolonged. Effort, however, is costly to the agent and unobservable to the principal. We study optimal dynamic contracts that always induce the agent to exert effort while maximizing the principal's profits. We formulate the contract design problem as a stochastic optimal control model with incentive constraints in continuous time over an infinite horizon. Although we consider the contract space that allows payments and potential contract termination time to take general forms, the optimal contracts demonstrate simple and intuitive structures, making them easy to describe and implement in practice.

Session 5.5 in Room M1 – Probabilistic Methods for Inference and Learning
Chair: Mohsen Bayati

Beyond Binary Inference: Contrast-Specific Propensity Scores for Heterogeneous Treatment Effect Estimation

S. Han, National University of Singapore, Singapore, shashahan@u.nus.edu
J. Goh, National University of Singapore, Singapore, joelgoh@nus.edu.sg
F. Meng, Department of Health Services & Outcomes Research, National Health Group, Singapore, Fanwen_MENG@nhg.com.sg
D. B. Rubin, Department of Statistics, Harvard University, USA, rubin@stat.harvard.edu

When there are three or more treatments, heterogeneous treatment effect estimation from observational data is a challenging task, even under the classical "strongly ignorable treatment assignment" assumption. On the one hand, using classical (binary) propensity score matching for pairs of treatments at a time does not guarantee transitivity in the treatment effect; on the other, matching on the full (vector-valued) multidimensional propensity score comes with scalability and other issues. We propose a new randomized approach for this task, which employs a generalization of the classical propensity score, which we call the *contrast-specific propensity score*. Although our approach does not circumvent the problem of scalability, we show that analytically (under certain conditions) and numerically (on simulated data) that it is more robust to model mis-specification compared to matching on the full multidimensional propensity score, and therefore a more promising approach for implementation in a real application. We apply our approach on a real dataset of diabetic patients in Singapore to generate heterogeneous medication recommendations for this patient population and to empirically corroborate our theoretical findings.

Adaptive Sequential Experiments with Unknown Information Flows

Y. Gur, Stanford University, USA, ygur@stanford.edu
A. Momeni, Stanford University, USA, amomenis@stanford.edu

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 multi-armed bandit (MAB) framework that has been studied and applied under strong assumptions on the information collection process: 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, and 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, and study how salient characteristics of the information impact policy design and achievable performance. 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 our approach by designing a adaptive algorithms for new product recommendations when additional consumers may arrive to product pages directly from external sources.

Global Convergence Guarantees for Policy Gradient Methods

J. Bhandari, Columbia Univeristy, jb3618@columbia.edu
D. Russo, Columbia, University, djr2174@gsb.columbia.edu

Policy gradients methods are perhaps the most widely used class of reinforcement learning algorithms. These methods apply to complex, poorly understood, control problems by performing stochastic gradient descent over a parameterized class of polices. Unfortunately, even for simple control problems solvable by classical techniques, policy gradient algorithms face non-convex optimization problems and are widely understood to converge only to local minima.
This work identifies structural properties – shared by finite MDPs and several classic control problems – which guarantee policy gradient methods convergence to globally optimal solutions despite the non-convex nature of the objective function. When these conditions are relaxed, our work guarantees convergence to a near-optimal performance, where the limiting error depends on a notion of the expressive capacity of the policy class. The analysis builds on standard theory of policy iteration. It seems to offer a clarifying perspective on a segment of the operations literature that studies online gradient algorithms for setting base-stock levels in inventory control and on recent work by Fazel, Ge, Kakade, and Mesbahi (2018) who establish global convergence of policy gradient methods in linear quadratic control problems through an intricate analysis of the relevant matrices.

A General Tail Bound for Matrix Estimation

N. Hamidi, Stanford University, USA, hamidi@stanford.edu
M. Bayati, Stanford University, USA, bayati@stanford.edu

In this talk, we study the trace regression problem for estimating a low-rank matrix. This problem encompasses a wide family of matrix estimation questions such as matrix completion and multi-task learning. We prove a new tail bound for this estimation problem under general sampling distributions.

Session 5.6 in Room M2 – Doubly Stochastic and State Dependent Queueing Models
Chair: Andrew Daw

Affine Point Processes: Refinements to Large-Time Asymptotics

X. Gao, Chinese University of Hong Kong, Hong Kong, xfgao@se.cuhk.edu.hk
L. Zhu, Florida Sate University, USA, zhu@math.fsu.edu

Affine point processes are a class of simple point processes with self- and mutually-exciting properties, and they have found applications in many areas. In this paper, we obtain large-time asymptotic expansions in large deviations and refined central limit theorem for affine point processes, using the framework of mod-φ convergence. Our results extend the large-time limit theorems in the work of Zhang et al. (2015). The resulting explicit approximations for large deviation probabilities and tail expectations can be used as an alternative to importance sampling Monte Carlo simulations. We carry out numerical studies to illustrate our results, and present applications in credit risk analysis.

Many-Server Queues with Autoregressive Inputs

X. Sun, Columbia University, USA, xs2235@columbia.edu

Recent studies have revealed the presence of significant autocorrelation and overdispersion in arrival data at large call centers. Motivated by these findings, we study a class of queueing systems where customers arrive according to a doubly stochastic Poisson point process whose intensities are driven by a time-dependent Cox-Ingersoll-Ross (CIR) process. The nonnegativity and autoregressive feature of the CIR process makes it a good candidate for modeling temporary dips and surges in arrivals. We conduct performance analysis of such systems. In particular, we study asymptotic performances such as the queue length and customer delays under suitable heavy traffic regimes. The results acknowledge the presence of autoregressive structure in arrivals and produce operational insights into staffing decisions.

Optimal Control Policies for an $M/M/1$ Queue with a Removable Server and Dynamic Service Rates

P. Badian-Pessot, Cornell University, USA, plb93@cornell.edu
D. Down, McMaster University, Canada, downd@mcmaster.ca
M. Lewis, Cornell University, USA, mark.lewis@cornell.edu

We consider an $M/M/1$ queue with a removable server that dynamically chooses its service rate from a set of finitely many rates. If the server is off, the system must warm up for a random, exponentially distributed amount of time, before it can begin processing jobs. We show under the average cost criterion, that work conserving policies are optimal. We then demonstrate the optimal policy can be characterized by a threshold for turning on the server and the optimal service rate increases monotonically with the number in system. Finally, we present some numerical experiments to provide insights into the practicality of having both a removable server and service rate control.

The Queue-Hawkes Process: Ephemeral Self-Excitement

A. Daw, Cornell University, United States, amd399@cornell.edu
J. Pender, Cornell University, United States, jjp274@cornell.edu

Across a wide variety of applications, the self-exciting Hawkes process has been used to model the history of events influencing future occurrences. In this talk, we define a novel generalization of the Hawkes process called the *Queue-Hawkes process*. This new stochastic process combines the dynamics of a self-exciting process and an infinite server queueing model: arrivals increase the arrival rate, but departures decrease it. By comparison to the Hawkes process, the Queue-Hawkes process is self-excitement on a system rather than on a sequence, making it an ephemerally self-exciting process. Our study of this model includes exploration of the process itself, investigation of relationships between self-exciting processes, and connections to well-known stochastic models such as branching processes, random walks, epidemics, and Bayesian mixture models. Our results for the Queue-Hawkes process include deriving a law of large numbers, fluid limits, and diffusion limit bounds for this new process. Furthermore, we prove a batch scaling construction of general Hawkes processes from a special affine case of the Queue-Hawkes process, which both provides insight into the Hawkes process and motivates the Affine Queue-Hawkes process as an attractive self-exciting process in its own right.

Session 5.7 in Rooms M5 and M6 – Stochastic Optimal Control in Finance
Chair: Duy-Minh Dang

Mean-Quadratic Variation (MQV) portfolio optimisation as an alternative to Time-consistent Mean-Variance (TCMV) optimisation

P. M. Van Staden, University of Queensland, Australia, pieter.vanstaden@gmail.com
D. M. Dang, University of Queensland, Australia, duyminh.dang@uq.edu.au
P. A. Forsyth, University of Waterloo, Canada, paforsyt@uwaterloo.ca

This talk investigates the Mean-Quadratic Variation (MQV) portfolio optimisation problem and its relationship to the Time-consistent Mean-Variance (TCMV) portfolio optimisation problem. We discuss conditions under which the two problems are (i) identical with respect to Mean-Variance trade-offs, and (ii) equivalent, i.e. have the same value function and optimal control. In order to compare the MQV and TCMV problems in a more realistic setting which involves investment constraints and modelling assumptions for which analytical solutions are not known to exist, we present an efficient partial integro-differential equation (PIDE) method for determining the optimal control for the MQV problem and discuss the associated convergence proof. We show that MQV investor achieves essentially the same results concerning terminal wealth as TCMV investor, but the MQV-optimal investment process has more desirable risk characteristics from the perspective of long-term investors with fixed investment time horizons. As a result, MQV portfolio optimisation is a potentially desirable alternative to TCMV optimisation.

Numerical methods for Guaranteed Minimum Withdrawal Benefits (GMWBs) as a continuous impulse control problem

Y. Lu, University of Queensland, Australia, yaowenlu@outlook.com
D. M. Dang, University of Queensland, Australia, duyminh.dang@uq.edu.au
P. A. Forsyth, University of Waterloo, Canada, paforsyt@uwaterloo.ca

We propose an ϵ-monotone Fourier method for solving a continuous impulse stochastic control formulation for valuation problem of Guaranteed Minimum Withdrawal Benefits (GMWBs) under jump diffusion models. We prove the convergence of our scheme to the viscosity solution of the continuous formulation, although the proposed method is not strictly monotone in the viscosity sense. Numerical experiments indicating the accuracy of the proposed method are presented.

An integration method for mean-variance portfolio optimisation under the jump-extended Heston model

H. Zhang, The University of Queensland, Australia, hanwen.zhang1@uq.net.au
D. Dang, The University of Queensland, Australia, duyminh.dang@uq.edu.au

We present an efficient epsilon-monotone numerical integration method for mean-variance (MV) portfolio optimisation problems under the jump-extended Heston model. This project aims to narrow research gaps in realistic modelling and efficient numerical methods for MV portfolio optimisation. The numerical scheme is guaranteed to be monotone within an infinitesimal tolerance epsilon, stable, and satisfy an epsilon-discrete comparison principle, and hence the convergence of the method can be attained. Numerical results, validated by Monte Carlo simulation, indicate that the proposed method is efficient.

Mean-risk portfolio optimisation

D. M. Dang, The University of Queensland, duyminh.dang@uq.edu.au

In this talk, we present survey recent interesting developments in mean-risk portfolio optimisation. Issues such as numerical methods, time(in)consistency of the controls , and robustness with respect to the model misspecifications, will be discussed.

This is a joint work with Peter Forsyth (Waterloo), Pieter M. van Staden, and Hanwen Zhang (University of Queensland)

Session 5.8 in Rooms M7 and M8 – Healthcare II
Chair: Ad Ridder

Modelling of reporting behaviour in the FluTracking surveillance system

D. Liu, The University of Adelaide, Australia, dennis.liu01@adelaide.edu.au
L. Mitchell, The University of Adelaide, Australia, lewis.mitchell@adelaide.edu.au
S. Carlson, Hunter New England Population Health, Australia, sandra.carlson@hnehealth.nsw.gov.au
J. V. Ross, The University of Adelaide, Australia, joshua.ross@adelaide.edu.au

Modelling the spread of influenza across Australia is of substantial public health concern. However, there are many challenges in verifying the accuracy of models, given that disease transmission in the general population is largely unobserved. Online participatory health surveillance systems attempt to address this challenge by providing a convenient and near real time platform for self-reporting of symptoms. FluTracking is one such platform for monitoring influenza-like-illness (ILI) in Australia, where participants are requested once a week to respond to an online survey, detailing any ILI symptoms experienced. Due to the voluntary nature of the platform, individuals reporting behaviour may vary over time. This leads to implications for modelling ILI incidence from the FluTracking system. For example, studies have shown users may be more likely to participate if they report on behalf of others, or report having ILI in their first report. Not considering this behaviour may lead to bias in estimates of ILI prevalence. In this work, we analyse weekly reports from the FluTracking system from May 2011 to October 2017 and use Bayesian logistic regression to model the probability of an individual reporting in a given week and estimate ILI prevalence in Australia for a given week. We then compare this to conventional naive estimates of ILI and show that consideration of the voluntary nature of this data should be considered when deriving estimates from the dataset.

Increasing the Health Benefits of In-home Chronic Care by Optimal Service Timing

M. Fackrell, The University of Melbourne, fackrell@unimelb.edu.au
J. Fu, The University of Melbourne, jing.fu@unimelb.edu.au
P. G. Taylor, The University of Melbourne, taylorpg@unimelb.edu.au
A. Tirdad, The University of Melbourne, ali.tirdad@unimelb.edu.au

We consider an in-home care service for patients with chronic conditions. Patients who have presented at an emergency department are assessed, and if deemed eligible, are invited to join the program. Once enrolled, patients are contacted frequently and asked a number of questions on how they feel about their current state of health. Depending on the answers they give, patients may be directed to an appropriate health service or admitted to hospital. The main goal of the service is to keep patients with chronic conditions healthier at home, in order to prevent, or at least reduce, avoidable hospitalisations and visits to the emergency department. One important consideration for in-home care service planners is to determine, each day, who to contact, and in which order of priority. We model this part of the program as a restless multi-armed bandit problem and aim at maximising the collective quality-adjusted-life-years of all the patients. By applying the idea of Whittle relaxation, we determine a tractable, priority-based algorithm which is demonstrated to outperform a myopic policy algorithm in our simulation results.

On congestion probabilities for ICU-SDU systems with overflow

N. M. van Dijk, University of Twente, The Netherlands, n.m.vandijk@utwente.nl
B. H. Schilstra, University of Twente, The Netherlands, b.h.schilstra@student.utwente.nl

Intensive care can be seen as one of the prime functions within hospitals. The availability of an ICU bed is thus of considerable interest. By integrating the intensive care unit (ICU) system with a so-called step down unit (SDU), the availability can be greatly enhanced. Closed form expressions are provided, as based on product form results, to show the effect of an SDU as well as to provide bounds on the congestion probability of an ICU in conjunction with an SDU. Numerical support is provided as based on realistic numbers. The results are of both theoretical and practical interest. Theoretical, as these cover a serial overflow with distinguishable patient types. Practical, as these provide bounds for dimensioning intensive care and step down capacities.

Minimizing bed occupancy variance by scheduling patients under uncertainty

A. Ridder, Vrije Universiteit Amsterdam, Netherlands, ad.ridder@vu.nl
A. van den Broek d'Obrenan, Vrije Universiteit Amsterdam, Netherlands
D. Roubos, HOTflo company, Netherlands
L. Stougie, Centrum voor Wiskunde en Informatica, Netherlands

In this paper we consider the problem of scheduling patients in allocated surgery blocks in a Master Surgical Schedule. Whereas extensive research efforts have been devoted on how to use the available surgery blocks as effectively as possible, little attention has been paid to the effect on the necessary bed capacity in hospital wards. However, for an efficient use of a hospital both types of resources should be taken into account. More specifically, large probabilities of overtime in each surgery block are undesirable and costly, while large fluctuations in the number of used beds requires extra buffer capacity and makes the staff planning more challenging. The stochastic nature of surgery durations and length of stay on a ward hinders the use of classical techniques. Transforming the stochastic problem into a deterministic problem does not result into practically feasible solutions. In this paper we develop a technique to solve the stochastic scheduling problem, whose primary objective it to minimize variation in the necessary bed capacity, while maximizing the number of patients operated, and minimizing the maximum waiting time, and guaranteeing a small probability of overtime in surgery blocks, and not exceeding the number of available beds at wards. The method starts with solving an Integer Linear Programming (ILP) formulation of the problem, and then simulation and local search techniques are applied to guarantee small probabilities of overtime and to improve upon the ILP solution. Numerical experiments applied to a Dutch hospital show promising results.

Thursday 15:00–16:30

Session 6.2 in Room P3 – Large Scale Interacting Queueing Networks
Chair: Ayalvadi Ganesh

Interference Queueing Networks on Grids

A. Sankararaman, The University of Texas at Austin, Austin TX, USA abishek@utexas.edu
F. Baccelli, The University of Texas at Austin, Austin TX, USA, baccelli@math.utexas.edu
S. Foss, Heriot-Watt University, Edinburgh, UK and Novosibirsk State University and Sobolev Institute of Mathematics, Novosibirsk, Russia, foss@hw.ac.uk

In this talk, we consider a countably infinite collection of interacting queues, with a queue located at each point of the d-dimensional integer grid. The queues have independent Poisson arrivals, but dependent service rates that are translation invariant in space. The service discipline of the processor sharing type, with the service rate in each queue slowed down, when the neighboring queues have a larger workload. The interactions are translation invariant in space and is neither of the Jackson Networks type,nor of the mean-field type. Coupling and percolation techniques are first used to show that this dynamics has well defined trajectories. Coupling from the past techniques are then proposed to build its minimal stationary regime. The rate conservation principle of Palm calculus is then used to identify the stability condition of this system, where the notion of stability is appropriately defined for an infinite dimensional process. We show that the identified condition is also necessary in certain special cases and conjecture it to be true in all cases. Remarkably, the rate conservation principle also provides a closed form expression for the mean queue size. When the stability condition holds, this minimal solution is the unique translation invariant stationary regime. In addition, there exists a range of small initial conditions for which the dynamics is attracted to the minimal regime. Nevertheless,there exists another range of larger though finite initial conditions for which the dynamics diverges, even though stability criterion holds.

Stability under utility-maximising allocations: from finite to infinite networks

S. Shneer, Heriot-Watt University, UK, V.Shneer@hw.ac.uk
A. Stolyar, University of Illinois at Urbana-Champaign, USA, stolyar@illinois.edu

We consider infinite networks of interacting queues governed by utility-maximising service-rate allocations, in both discrete and continuous time. We show how moment bounds for stationary distributions of finite networks may be used to prove stability of infinite networks. We also prove stability of some symmetric multi-hop infinite networks.

Phase transitions in a generalized hard core model (k - casting)

K. Ramanan, Brown University, USA, Kavita_Ramanan@brown.edu
I. Ziedins, University of Auckland, New Zealand, i.ziedins@auckland.ac.nz

We consider a loss network with the structure of a tree network. Each internal node of the network has d edges attached to it, with each edge having unit capacity. An arrival requests unit capacity on $k \leq d$ edges out of a node. Arrivals to each possible route occur as a Poisson process of rate ν. Call durations are generally distributed with mean 1. In the case $k = d$ it is well known that phase transitions can occur if $k > 2$. Here we consider the case $k < d$. We show that the model has a unique invariant Gibbs measure for all ν. Furthermore, if $k = 2$, then no phase transition occurs. For $k > 2$, however, there is a complex interaction between k and d which determines whether a phase transition occurs or not.

Large deviations for Cox processes and $Cox/G/\infty$ queues, with a biological application

J. Dean, Harvard University, justindean9208@gmail.com
A. Ganesh, University of Bristol, UK, a.ganesh@bristol.ac.uk
E. Crane, University of Bristol, UK, edward.crane@bristol.ac.uk

We show that a sequence of Cox processes on a Polish space E satisfy a large deviation principle (LDP), provided their directing measures do so on the space of finite measures on E equipped with the weak topology. Next, we consider a sequence of infinite server queue with general iid service times, where the arrivals constitute Cox processes with translation invariant directing measures assumed to satisfy an LDP. We show that the corresponding sequence of queue occupancy measures also satisfy an LDP. These problems were motivated by the problem of describing fluctuations of molecule numbers in biochemical reaction networks within cells.

Session 6.3 in Room P4 – Stochastic Models and Matching Queues
Chair: Pascal Moyal

Dynamic Matching for Real-Time Ridesharing

A. R. Ward, The University of Chicago, USA, amy.ward@chicagobooth.edu
E. Özkan, Koç University, Turkey, erhozkan@koc.edu.tr

In a ridesharing system, arriving customers must be matched with available drivers. These decisions affect the overall number of customers matched, because they impact whether or not future available drivers will be close to the locations of arriving customers. A common policy used in practice is the closest driver (CD) policy that offers an arriving customer the closest driver. This is an attractive policy because it is simple and easy to implement. However, we expect that parameter-based policies can achieve better performance.
We propose matching policies based on a continuous linear program (CLP) that accounts for (i) the differing arrival rates of customers and drivers in different areas of the city, (ii) how long customers are willing to wait for driver pick-up, (iii) how long drivers are willing to wait for a customer, and (vi) the time-varying nature of all the aforementioned parameters. We prove asymptotic optimality of a forward-looking CLP-based policy in a large market regime and of a myopic LP-based matching policy when drivers are fully utilized. When pricing affects customer and driver arrival rates, and parameters are time homogeneous, we show that asymptotically optimal joint pricing and matching decisions lead to fully utilized drivers under mild conditions.

A token-based central queue with order-independent service rates

U. Ayesta, CNRS-IRIT/Université de Toulouse, France and IKERBASQUE - Basque Foundation for Science/University of the Basque Country, Spain, urtzi.ayesta@irit.fr
T. Bodas, CNRS-LAAS/Université de Toulouse, France, tejaspbodas@gmail.com
J. L. Dorsman, University of Amsterdam, The Netherlands, j.l.dorsman@uva.nl
I. M. Verloop, CNRS-IRIT/Université de Toulouse, France, maaike.verloop@enseeiht.fr

In this presentation, we regard a token-based central queue with multiple customer types, where customers of each type have an associated set of compatible tokens. Customers may only receive service when they have claimed a compatible token. If upon arrival more than one compatible token is available, an *assignment rule* determines which token will be claimed. When, however, there is no compatible token available at all, the customer waits in the central queue. An important feature of this queue is that the *service rate* obtained by a customer is state-dependent, i.e. it depends on which of the tokens are claimed and on the number of customers present in the system.
We will see in this presentation that, provided the assignment rule and the service rates satisfy certain conditions, the steady-state distribution of the general token-based central queue has a product form. In particular, we will see that our model subsumes known families of queueing models that have product-form steady-state distributions, including the order-independent queue described by Krzesinski (2011) and the multi-type customer and server queueing system studied by Visschers, Adan, and Weiss (2012). The product form of the steady-state distribution allows for tractable expressions for relevant performance measures such as the sojourn time of customers and the number of customers present in the system. These expressions can be used to analyse many relevant models, such as stochastic matching models and multi-server models with redundancy.

FCFS Dynamic Matching Models

G. Weiss, The University of Haifa, Israel, gweiss@stat.haifa.ac.il

Parallel service systems have several types of customers (participants, recipients), and several types of servers (agents, donors) that arrive along time and are matched according to a compatibility graph, with a focus on First come first served (FCFS) matching. Applications to call centers, organ transplants, auctions and markets, occur widely. As a queuing system this is generally a highly intractable model, but under Poisson exponential assumption it has partial balance and product form behavior. This leads to the simplified abstraction of FCFS matching of i.i.d. types subject to compatibility graph, where a very complete theory can be obtained. I will discuss this and some open problems that arise with many server scaling.

Coupling from the past of stochastic matching models

P. Moyal, Université de Lorraine, France, pascal.moyal@univ-lorraine.fri
A. Busic, INRIA/ENS, France, ana.busic@inria.fr
J. Mairesse, CNRS/UPMC, France, jean.mairesse@lip6.fr

We address stochastic matching models in the following sense: items enter a system at random times, and belong to designated classes. The compatibility between classe are given by a matching graph. The items use the system as an interface to be matched by pairs, and matchings are allowed only between items of compatible classes. A matching policy determines the matching to be executed in cases of multiple choices.
After briefly recalling general stability conditions in function of the geometry of the matching graph and of the matching policy, we present the construction of the stationary state of stable systems, *via* a coupling-from-the past scheme *à la* Loynes. This backwards scheme, which is the crucial first step toward perfect simulation, mainly relies on a block-wise sub-additivity property that is satisfied by most usual matching policies.
We address in this fashion, two sub-classes of this general class of models: general matching models (general matching graph and single arrivals) and the so-called extended bipartite models (bipartite matching graph and pairwise arrivals).

Session 6.4 in Room P5 – Approximations of Stochastic Processes
Chair: Giang Nguyen & Oscar Peralta

The Asymptotic Behavior of the Time-Varying Supermarket Model

H. M. Jansen, The University of Queensland, Australia, h.m.jansen@uq.edu.au

In the supermarket model, there are n parallel servers, each having its own queue. Arriving customers select d queues at random and join the shortest among them. Much is known about this model if its parameters are constant, but it is not clear how it behaves if the arrival rate varies throughout the day, for instance. In this talk, we take a closer look at the supermarket model with time-varying parameters under the Join-the-Shortest-Queue policy (the special case in which $d = n$). We explore in particular how the time-varying nature of the parameters influences this model on the diffusion scale as the number of servers becomes large.

Simulation-based assessment of the stationary tail distribution of a stochastic differential equation

K. Bisewski, CWI, The Netherlands, K.L.Bisewski@cwi.nl
D. Crommelin, CWI, The Netherlands, Daan.Crommelin@cwi.nl
M. Mandjes, University of Amsterdam, The Netherlands, m.r.h.mandjes@uva.nl

A commonly used approach to analyzing stochastic differential equations (SDEs) relies on performing Monte Carlo simulation with a discrete-time counterpart. In this talk I'll discuss the impact of such a time-discretization when assessing the stationary tail distribution. For a family of semi-implicit Euler discretization schemes with time-step $h > 0$, we quantify the relative error due to the discretization, as a function of h and the exceedance level x. By studying the existence of certain (polynomial and exponential) moments, using a sequence of prototypical examples, we demonstrate that this error may tend to 0 or infinity. The results show that the original shape of the tail can be heavily affected by the discretization. The cases studied indicate that one has to be very careful when estimating the stationary tail distribution using Euler discretization schemes.

Rate of strong convergence of stochastic fluid processes to Markov-modulated Brownian motion

G. Nguyen, The University of Adelaide, Australia, giang.nguyen@adelaide.edu.au
O. Peralta, The University of Adelaide, Australia, oscar.peraltagutierrez@adelaide.edu.au

In Latouche and Nguyen (2015), the authors constructed a sequence of stochastic fluid processes and showed that it converges weakly to a Markov-modulated Brownian motion $\{(X_t, \varphi_t)\}_{t\geq 0}$.
Here, we construct another sequence of stochastic fluid processes, with different characteristics to the ones considered in Latouche and Nguyen (2015), and show that it converges strongly to $\{(X_t, \varphi_t)\}_{t\geq 0}$. We also show that the rate of this almost sure convergence is proportional to $n^{-1/2}\log n$.
When reduced to the special case of standard Brownian motion, our convergence rate is an improvement over that obtained by Gorostiza and Griego (1980), which is proportional to $n^{-1/2}(\log n)^{5/2}$.

Convergence of a bivariate flip–flop process

G. Latouche, Université Libre de Bruxelles, Belgium, latouche@ulb.ac.be
G. Nguyen, The University of Adelaide, Australia, giang.nguyen@adelaide.edu.au
O. Peralta, The University of Adelaide, Australia, oscar.peraltagutierrez@adelaide.edu.au

Flip–flop processes refer to a family of stochastic fluid processes which converge to either a standard Brownian motion (SBM) or to a Markov modulated Brownian motion (MMBM). In recent years it has been shown that complex distributional aspects of the univariate SBM and MMBM can be studied through the limiting behaviour of flip–flop processes. In this talk, we define a class of bivariate flip–flop processes whose marginals converge to SBMs and are dependent on each other. We exhibit an example in which the limiting bivariate process is not Gaussian, nonetheless, it possesses desirable qualities, such as being tractable and having a time–varying correlation coefficient function.

Session 6.5 in Room M1 – Learning in Sequential Decision Problems
Chair: Dan Russo

Contextual Bandits with a Low-Rank Structure

N. Hamidi, Stanford University, USA, hamidi@stanford.edu
M. Bayati, Stanford University, USA, bayati@stanford.edu
K. Gupta, Airbnb, USA, guptakapil@gmail.com

We consider the k-armed stochastic contextual bandit problem with d dimensional features, when both k and d can be large. To the best of our knowledge, all existing algorithm for this problem have a regret bound that scale as polynomials of degree at least two in k and d. In this talk we introduce and theoretically analyze a new algorithm with a regret that scales $k + d$ when the matrix of unknown parameters has a low-rank structure.

Delay-Predictability tradeoffs in reaching a secret goal

J. N. Tsitsiklis LIDS, Massachusetts Institute of Technology, jnt@mit.edu
K. Xu Graduate School of Business, Stanford University, kuangxu@stanford.edu

We formulate a model of sequential decision-making, dubbed the Goal Prediction game, to study the extent to which an overseeing adversary can predict the final goal of an agent who tries to reach that goal quickly, through a sequence of intermediate actions. Our formulation is motivated by the increasing ubiquity of large-scale surveillance and data collection infrastructures, which can be used to predict an agent's intentions and future actions, despite the agent's desire for privacy.
Our main result shows that with a carefully chosen strategy, the predictability of the agent's goal can be made inversely proportional to the time she is willing to spend in reaching it, and that this is essentially the best possible. Moreover, this characterization depends on the topology of the agent's state space only through its diameter.

AlphaGo Zero, Monte Carlo Tree Search and Self-Play: Towards Theoretical Foundations

D. Shah, MIT, devavrat@mit.edu
Q. Xie, Cornell, qiaomin.xie@cornell.edu
Z. Xu, MIT, zhixu@mit.edu

AlphaGo Zero (AGZ) by Silver et al. (2017) introduced a new tabula rasa reinforcement learning algorithm that has achieved superhuman performance in the games of Go, Chess, and Shogi with no prior knowledge other than the rules of the game. Methodically, there are two key innovations: (a) use of Monte-Carlo Tree Search (MCTS) with Supervised Learning (SL) through well-engineered Neural Network for learning the policy, and (b) use of self-play for generating training examples. The remarkable empirical success naturally raises the following question: can we explain this empirical success theoretically, or is it simply ingenious engineering?
In this talk, we shall argue that MCTS with expressive enough SL learns optimal policy at nearly minimax optimal rate. In the process, we correct a fundamental error in the well cited MCTS method and its proof. Interestingly enough, the AGZ had already utilized this correction in its implementation. However, our theoretical analysis suggests further improvement over AGZ's implementation. Our results hold for both traditional setting of infinite horizon Markov Decision Process (MDP) as well as the setting of self-play modeled as Robust MDP. Beyond two player games where AGZ has been effective, we show empirical success for a representative application from network scheduling.

Costly Learning Manipulation: Belief Distortion Through Information Dissemination

J. Keppo, NUS Business School, Singapore, keppo@nus.edu.sg
M. Kim, UBC Sauder School of Business, Canada, mike.kim@sauder.ubc.ca
X. Zhang, UBC Sauder School of Business, Canada, xinyuan.zhang@sauder.ubc.ca

We consider optimal manipulation of an uninformed learner through strategic design of data provision. The problem is motivated by communication settings in which a credible information provider can misrepresent the information sent to the learner, whose resistance to the manipulation creates an economic cost to the information provider. Applications include advertising management and political propaganda. Specifically, the learner is assumed Bayesian, and the information provider specifies both the precision and accuracy of the information disseminated to the learner, aiming to distort his posterior beliefs hereafter. We show that the dynamic control problem of interest has a tractable reformulation, and we explicitly characterize, in closed form, the optimal manipulation policy. We establish that value overstatement by the information provider optimally arises, and the manipulation is either lumpy or smooth, depending on how the cost varies over time. The explicit solution also allows us to fully describe the time trajectory of the learner's belief process under distortion, rendering the manipulation dynamics perspicuous.

Session 6.6 in Room M2 – Diffusions and Simulation
Chair: Vlad Margarint

Exact Simulation of Coupled Wright–Fisher Diffusions

C. García-Pareja, Karolinska Institutet, Sweden, celia.garcia.pareja@ki.se
H. Hult, KTH Royal Institute of Technology, Sweden, hult@kth.se
T. Koski, KTH Royal Institute of Technology, Sweden, tjtkoski@kth.se

In this talk we introduce an exact rejection algorithm for simulating paths of a family of multivariate Wright–Fisher diffusions, the coupled Wright–Fisher diffusion, which models the co-evolution of multiple genetic traits at different locations on the genome. Our algorithm uses independent one-dimensional neutral Wright–Fisher diffusions as candidate proposals. The candidates can be sampled exactly by means of existing algorithms and are only needed at a finite number of points. Once a candidate is accepted, the remaining of the path can be recovered by sampling from a neutral multivariate Wright–Fisher bridge, for which we also provide an exact sampling strategy. The technique relies on a modification of the alternating series method and extends existing algorithms that are currently only available for the one-dimensional case. Finally, our algorithm's complexity is derived and its performance evaluated in a simulation study which shows promising results.

Unbiased Simulation of Multivariate Jump-Diffusions

G. Chen, Stanford University, USA, guanting@stanford.edu
K. Giesecke, Stanford University, USA, giesecke@stanford.edu
A. Shkolnik, University of California, Santa Barbara, USA, shkolnik@ucsb.edu

Techniques for the simulation of stochastic differential equations (SDEs) have attracted a significant amount of interest in the Monte Carlo methods and applied probability communities. Recent breakthroughs on generic algorithms for multivariate diffusions include exact schemes and unbiased methods. However, the addition of state-dependent jumps to the SDE presents significant challenges for these approaches. We develop the first (to our knowledge) method for simulation of multivariate jump- diffusions with general drift, volatility and jump-intensity coefficients. The approach allows one to embed essentially any method for simulating the pure diffusion in between the jump times. While exact sampling of state-dependent jumps is in general not feasible, we propose a change of probability measure to circumvent this obstacle. The change of measure is induced by certain point process martingales and is of independent interest. Numerical results for applications in finance demonstrate the advantages of the algorithm.

Pathwise and probabilistic analysis in the context of SLE

V. Margarint, University of Oxford, UK, margarint@maths.ox.ac.uk

In this talk, I am going to report on some research projects at the interface between Rough Paths Theory and Schramm–Loewner evolution (SLE). In the study of Rough Differential Equations questions such as uniqueness/non uniqueness of solutions depending on the behaviour of parameters of the equation appear naturally. We adapt these type of questions to the study of the backward Loewner Differential Equation in the upper half-plane. The main ideas concern the restart of the backward Loewner Differential Equation from the singularity. I will cover two parts, in which I am going to describe some general tools that lead to a better understanding of the dynamics in the closed upper half plane under the backward Loewner flow driven by Brownian Motion. In the first part, I will describe a coordinate change of the Loewner Equation in which we obtain in a random time change, a stochastic dynamics on a specific line in $\mathbb{H}$, that is depending only on the argument of the points. In this section, I am going to cover an analysis of this dynamics and related objects. We will further use the aforementioned estimates to reveal new information about the SLE trace at fixed capacity parametrisation times. In the second part, I will explain how the uniqueness/non-uniqueness of solutions of the backward Loewner Equation driven by Brownian Motion, started from the singularity, can be understood using some property of the dynamics on the boundary. The main tool is the characterisation of boundary behaviour of Bessel processes of low and negative dimensions. In this new language, we will recover some new information about the SLE traces.
If time permits, I am going to discuss also a new study concerning Euler approximations of solutions for singular Rough Differential Equations started from the singularity with an application to Loewner Differential Equation. This last project opens a new chapter in Rough Path Theory.

Session 6.7 in Rooms M5 and M6 – Monte Carlo Methods
Chair: Pierre L'Ecuyer

Control variates for censored Monte Carlo simulation

E. Akin, Naval Postgraduate School, USA, ewakin@nps.edu
R. Szechtman, Naval Postgraduate School, USA, rszechtm@nps.edu

We consider the problem of estimating the mean of a random variable X via Monte Carlo simulation, in the presence of censoring. In such setting, the value of each random sample X_i, $i = 1, \ldots, n$, is known if $X_i \in A_i$, where the ranges A_i are also random. We propose a control variates estimator that uses the information about the censored observations, and draw connections with the empirical likelihood estimator for the mean of X.

Multilevel Monte-Carlo method and lower/upper bounds in Initial Margin computations

F. Bourgey, Ecole Polytechnique, France, florian.bourgey@polytechnique.edu
S. De Marco, Ecole Polytechnique, France, stefano.de-marco@polytechnique.edu
E. Gobet, Ecole Polytechnique, France, emmanuel.gobet@polytechnique.edu
A. Zhou, Université Paris-Est, CERMICS (ENPC), France, alexandre.zhou@enpc.fr

The Multi-level Monte-Carlo (MLMC) method developed by Giles (2008) has been successfully applied in innumerous fields of stochastic simulation. Quoting Giles (2015), MLMC "reduces the computational cost [*with respect to standard Monte-Carlo*] by performing most simulations with low accuracy at a correspondingly low cost, with relatively few simulations being performed at high accuracy and a high cost". A natural application of this method is the evaluation of *nested* expectation of the form $E\big[g(E[f(X,Y)|X])\big]$, where f, g are given functions and (X, Y) a couple of independent random variables. Apart from the pricing of American-type derivatives, such computations arise in a large variety of risk valuations (VaR or CVaR of a portfolio, CVA), or in the assessment of margin costs of centrally cleared portfolios. In this work, we focus on the computation of Initial margins. We analyze the properties and asymptotically optimal choices of MLMC estimators in practical situations of limited regularity of the outer function g (with singularities in the first derivative). In parallel, we investigate upper and lower bounds for nested expectations as above, in the spirit of primal/dual algorithms for stochastic control problems.

Quasi-Monte Carlo Density Estimation

A. Ben Abdellah, Université de Montréal, Canada, amalbenabdellah@gmail.com
P. L'Ecuyer, Université de Montréal, Canada, lecuyer@iro.umontreal.ca
A. B. Owen, Stanford University, USA, owen@stanford.edu
F. Puchhammer, Université de Montréal, Canada, florian.puchhammer@umontreal.ca

We consider the problem of estimating the density of a random variable X that can be sampled exactly by Monte Carlo (MC). We study the effectiveness of replacing the underlying independent uniform random numbers by randomized quasi Monte Carlo (RQMC) points or by stratified random numbers, to reduce the integrated variance (IV) and the mean integrated square error (MISE) for kernel density estimators. We provide theoretical upper bounds on the IV and MISE with these sampling methods, and give conditions under which these bounds converge at a faster rate than the MC rates. The bias remains the same as for MC. We find that the IV bounds obtained via a traditional Koksma–Hlawka-type inequality for RQMC are too loose to be useful when the dimension of the problem exceeds a few units. We describe an alternative way to estimate the IV, a good bandwidth, and the MISE, under RQMC or stratification, and we show empirically that the IV and MISE can be reduced significantly by RQMC even for some high-dimensional problems.

Session 6.8 in Rooms M7 and M8 – Applications in Economics
Chair: Nick Arnosti

Misconducts in Organizations

M. Zhou, National University of Singapore, Singapore, minglong_zhou@u.nus.edu
J. Keppo, National University of Singapore, Singapore, keppo@nus.edu.sg
E. Jokivuolle, Bank of Finland, Finland, Esa.Jokivuolle@bof.fi

Insider misconducts has caught public attention during the last few years. For instance, in 2017, Bupa's insider threat data leak affected over 500,000 international health insurance customers. These misconducts arise partly from high-powered financial incentives that are influenced by laws, regulations, and firms' own policies. In this paper, we model the interaction among lawmakers, firms, and their employees under the risk of insider misconducts. Agents (employees), with private information on their type and action, are potential cheaters. Principals (firms) interact with agents of uncertain types, and enforce a certain level of costly monitoring to maximize the firm value. Regulators (lawmakers) affect the agents and principals' incentives by imposing legal penalties. First, we solve a simultaneous game with asymmetric information to explain employee-level misconducts in firms. Our model supports empirical evidence that middle management typically causes large-scale misconducts. Second, we study how the equilibrium of the simultaneous game help in designing effective laws and regulations. Our results provide insights on why regulations differ in different countries and states.

Modelling Double Auctions with dynamic supply and Demand

B. Niknami, The University of Melbourne, Australia, bniknami@student.unimelb.edu.au
P. G. Taylor, The University of Melbourne, Australia, taylorpg@unimelb.edu.au

Double auctions are systems used in many financial markets to facilitate the exchange of assets according to predetermined regulations. As such, their structures play important roles in price formation and market health. A rigorous understanding of their behaviour would support informed decision making on the part of investors and regulators.
There have been some attempts to study simplified versions of these systems in the literature, where modest assumptions are made about the structure of the market, leading to interesting results. Some of these assumptions are that market participants submit orders independently of each other and the supply and demand functions are static in time. To make analysis feasible they also focus on fewer types of orders, often only looking at auctions with limit orders flowing into them. These facilitate the study of equilibrium behaviours for such systems. However, these conjectures are too restrictive for real life; but there have been attempts to generalise the above model in different manners.
We also aim to loosen the above assumptions, initially, to describe cases where the supply and demand functions are not constant. This will pave the way for a more general understanding of double auctions, which could promote future studies of the effects of regulation and trader behaviour on the market.

Optimal Market Thickness and Clearing

S. Loertscher, University of Melbourne, Australia, simonl@unimelb.edu.au
E. V. Muir, Stanford University, United States, evmuir@stanford.edu
P. G. Taylor, University of Melbourne, Australia, taylorpg@unimelb.edu.au

To determine the optimal market clearing policy that balances the benefits of thickness against costly delay, we solve a dynamic mechanism design model with sequential arrival of buyers and sellers with private information about their types. With discrete (binary) types, there is an efficient mechanism that is incentive compatible, individually rational and always balances the budget if the discount factor is sufficiently large (if storing at least one trade is efficient). For large discount factors, most welfare gains from dynamic mechanisms relative to instantaneous clearing are reaped by the simplest dynamic mechanism that clears at an optimally chosen, fixed frequency.

A Continuum Model of Stable Matching with Finite Capacities

N. Arnosti, Columbia Business School, nicholas.arnosti@gmail.com

We introduce a continuum model of stable matching. Unlike previous work, our model allows for arbitrary capacities, preferences, and priorities. We show there always exists a unique stable outcome, and that this outcome can be computed by repeatedly applying a monotone operator.
The model reproduces numerous simulation outcomes from the literature. For example, it accurately predicts the distribution of cutoffs in a many-to-one market with correlated preferences. Although the model features a continuum of agents, it gives nearly exact answers even for markets with ten or twenty participants on each side.
We close by discussing several applications. First, we use the model to analyze a game of strategic list formation, and discuss how this could be used by econometricians interested in estimating counterfactuals. Second, we use the model to study pricing questions in online marketplaces, and lottery design questions for school choice.

Friday 10:30–12:00

Session 7.2 in Room P3 – Resource Allocation for Emerging Applications
Chair: Siva Theja Maguluri

Transform Methods for Heavy-Traffic Analysis

D. Hurtado-Lange, Georgia Institute of Technology, United States of America, d.hurtado@gatech.edu
S. T. Maguluri, Georgia Institute of Technology, United States of America, siva.theja@gatech.edu

The drift method was recently developed as an alternative to the diffusion limit approach to study queueing systems. It was successfully used to obtain bounds on the moments of the scaled queue lengths, that are asymptotically tight in heavy-traffic, in a wide variety of systems including generalized switches, input-queued switches, bandwidth sharing networks, etc. In general, these moments are obtained by setting the drift of polynomial test functions to zero. Under certain conditions, this allows one to obtain the joint distribution of queue lengths. In this talk we present the Moment Generating Function (MGF) method, that simplifies the drift method, and provides a new perspective, unifying it with other transform based methods in the literature. We develop this method for queueing systems that satisfy the so called Complete Resource Pooling condition, such as load balancing systems under the Join the Shortest Queue (JSQ) and power-of-two choices policies, ad hoc wireless networks with interference constraints and fading.

State Dependent Control of Closed Queueing Networks with Application to Ride-Hailing

S. Banerjee, Cornell University, USA, sbanerjee@cornell.edu
Y. Kanoria, Columbia University, USA, ykanoria@columbia.edu
P. Qian, Columbia University, USA, pq2124@columbia.edu

Inspired by ride-hailing and bike-sharing systems, we study the design of state-dependent controls for a closed queueing network model. We focus on the assignment policy, where the platform can choose which nearby vehicle to assign to an incoming customer; if no units are available nearby, the request is dropped. The vehicle becomes available at the destination after dropping the customer. We study how to minimize the proportion of dropped requests in steady state.
We propose a family of simple state-dependent policies called Scaled MaxWeight (SMW) policies that dynamically manage the geographical distribution of supply. We prove that under the complete resource pooling (CRP) condition (analogous to the condition in Hall's marriage theorem), each SMW policy leads to exponential decay of demand-dropping probability as the number of supply units scales to infinity. Further, there is an SMW policy that achieves the **optimal** exponent among all assignment policies, and we analytically specify this policy in terms of the customer arrival rates for all source-destination pairs. The optimal SMW policy maintains high supply levels near structurally under-supplied locations. We also propose data-driven approaches for designing SMW policies and demonstrate excellent performance in simulations based on the NYC taxi dataset.

Delay and stability in distributed service systems with redundancy and random slowdowns

D. Gamarnik, MIT, Cambridge, USA, gamarnik@mit.edu
J. N. Tsitsiklis, MIT, Cambridge, USA, jnt@mit.edu
M. Zubeldia, MIT, Cambridge, USA, mar.zubeldia@gmail.com

We consider a large distributed service system consisting of n servers with infinite capacity FIFO queues, where each server can experience random slowdowns in its processing rate. Each incoming job consists of k identical tasks that can be executed in parallel, and that can be encoded into any number of at least k "replicas" of the same size (by introducing redundancy) so that the job is considered to be completed when *any* k replicas associated with it finish their service.
Our objective is to understand the best possible performance of such systems (in terms of stability region and delay) and to propose optimal policies, with emphasis on the asymptotic regime when n is large. In particular, we consider a broad family of control policies, which includes most policies considered in the literature, and work towards characterizing the fundamental tradeoff between the delay performance and the stability region, under two different models for the slowdowns: one where the slowdowns are independent of the job sizes, and a new one where slowdowns are job size dependent.

On the design of service systems when servers are strategic

R. Gopalakrishnan, Cornell University, Ithaca, New York, USA, rg584@cornell.edu

Many service systems are staffed by human servers who enjoy some discretion over their service times; however, queueing models with strategic service rates (motivated by systemic incentives) are only beginning to be studied. Recent work in this emerging area has shown that system policies (such as routing and staffing policies) that are provably optimal under classical models could end up being far from optimal as a result of undesirable strategic incentives created by these policies. A popular approach to capture this effect is to augment classical queueing models with a systemic utility model for strategic servers, according to which a server endogenously chooses their service rate to maximize their utility. Subsequently, the performance of the system is studied under a Nash equilibrium (if/when it exists) of the induced “server game”.
In this talk, I adopt a stylized utility model in which servers enjoy idle time, but incur a cost of effort; a joint optimization of the *routing policy* (over a large class of rate-based policies, e.g., Fastest/Slowest-Server-First) and the *system configuration* (pooled with a single queue vs. dedicated with parallel queues) is then studied. Under either configuration, the optimal routing policies outperform uniform random routing in terms of the mean response/waiting times. However, such system-optimal policies also minimize servers’ utilities, creating a moral dilemma for a system manager who also cares about employee satisfaction. Another interesting property of system-optimal policies is that servers work faster when there are more servers in the system, which adds a nontrivial angle to the classical question of optimal staffing. Finally, servers work faster under the dedicated configuration, but not necessarily enough to overcome its inherent systemic inefficiency. These results align qualitatively with observations from real-world human service systems.

Session 7.3 in Room P4 – Approximations and Controls for Queues
Chair: Yunan Liu

A Fluid Limit for an Overloaded Multiclass Many-Server Queue with General Reneging Distribution

A. Puha, California State University San Marcos, USA, apuha@csusm.edu
A. R. Ward, The University of Chicago, USA, amy.ward@chicagobooth.edu

We study scheduling in a many-server queue with general reneging distribution ($G/GI/N+GI$) and multiple customer classes. To do this, we specify a class of admissible control policies (rules for determining when to serve a given customer class) and formulate a fluid model. We establish a tightness result for sequences of fluid scaled admissible control policies satisfying mild conditions, and show that limit points of such sequences satisfy the set of fluid model equations. Then, we characterize the invariant states of the fluid model, and introduce a set of control policies, called Weighted Random Buffer Selection (WRBS), that capture the entire spectrum of invariant states. We prove that a suitably rescaled state descriptor for the queue operating under a specified WRBS policy converges to the unique fluid model solution for that WRBS policy.

Admission control for double-ended queueing systems in heavy traffic

X. Liu, Clemson University, USA, xliu9@clemson.edu
A. Weerasinghe, Iowa State University, USA, ananda@iastate.edu

Motivated from the study of trading systems, we consider a double-ended queue with one side of the queue holding sellers, and the other side for buyers. We assume that both sellers and buyers are impatient. Our goal is to design optimal capacity sizes for both sides of the queue to minimize an infinite horizon discounted cost functional, which consists of linear holding costs for waiting customers, and penalty costs for blocked customers, and customers who abandon. Note that customers arriving at a full queue will be blocked and customers who wait too long will abandon. Thus our control problem studies the tradeoff between blocking and abandonment. We study the queueing system in heavy traffic, and formulate a diffusion control problem (DCP). Solving the DCP, we develop asymptotically optimal capacity sizes for the system.

Staffing and Scheduling to Differentiate Service in Multiclass Time-Varying Service Systems

Y. Liu, North Carolina State University, NC, yliu48@ncsu.edu
X. Sun, University of Florida, FL, xs2235@columbia.edu
K. Hovey, North Carolina State University, NC, kahovey@ncsu.edu

Motivated by call centers and health care systems, we study a queueing system having multiple customer classes, nonstationary customer arrivals, and customer abandonment. We develop effective staffing rules (number of servers) and scheduling policies (assigning newly idle servers to a waiting customer from one of the classes), with the objective of achieving differentiated service levels for each customer class. One notable motivation of this research is the Canadian triage and acuity scale (CTAS) guideline that classifies patients in the emergency department (ED) into five acuity levels. In particular, CTAS requires that "level i patients need to be seen by a physician within w_i minutes $100\alpha_i\%$ of the time", with $(w_1, w_2, w_3, w_4, w_5) = (0, 15, 30, 60, 120)$ minutes, and $(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5) = (0.98, 0.95, 0.9, 0.85, 0.8)$. Our goal is to devise new control principles (via staffing and scheduling) to satisfy a constraint on the class-dependent tail probability of delay (TPoD): $P(W_i(t) > w_i) \leq \alpha_i$, that is, the probability that a class-i customer waits more than $w_i > 0$ does not exceed $\alpha_i \in (0, 1)$ at all times for all classes. Our new joint staffing and scheduling policy is both time dependent (which copes with the time variability in arrival pattern) and state dependent (which dynamically captures the stochastic variability in service times and arrival times). Effectiveness of our policy is substantiated by a many-server asymptotic optimality theorem in the efficiency-driven regime (our limit theorem shows that the TPoD $P(W_i(t) > w_i) \to \alpha_i$ when the system scale increases). We also conduct computer simulation experiments to provide engineering confirmation and to gain insights.

Optimal Repositioning of Urban Electric Vehicle Sharing Systems

X. Sun, Columbia University, USA, xs2235@columbia.edu

We consider a critically-loaded queueing network model of shared electric vehicles (EVs) traveling in an urban environment. We focus on reposition of empty EVs by which we dynamically control the flow of EVs in the network to minimized costs associated with control and customer losses. We formulate an optimal rate control problem for such a network with an ergodic cost criterion and show that the value function (i.e., minimum value of the cost) of the rate control problem for the network converges, under a suitable heavy traffic scaling, to that of a control problem for a set of controlled reflected diffusions. We characterize the optimal control policy for the limiting diffusion model via a viscosity solution to a Hamilton-Jacobi-Bellman equation. Furthermore, we show that the optimal control policy for the limit diffusion model can be used to construct asymptotically optimal control policies for the original network.

Session 7.4 in Room P5 – Branching Random Walks
Chair: Sophie Hautphenne & Peter Braunsteins

The global critical value of branching random walks: weak but hard!

F. Zucca, Politecnico di Milano, Italy, fabio.zucca@polimi.it
D. Bertacchi, Università di Milano-Bicocca, Italy, daniela.bertacchi@unimib.it

Given a continuous-time branching random walk, two thresholds can be identified: the local and global critical values. The first one has a fairly simple explicit characterization in terms of the first-moment matrix of the process. For the global one, there are a few possible characterizations based on functional inequalities; nevertheless, many explicit general expressions have been proposed in the last decade and they have all been rejected. We see where these expressions fail and under which conditions they hold.

The cardinality of extinction probability vectors of a branching random walk

D. Bertacchi, University of Milano-Bicocca, Italy, daniela.bertacchi@unimib.it
F. Zucca, Politecnico di Milano, Italy, fabio.zucca@polimi.it

Given a branching random walk on a set X, we study its extinction probability vectors, whose x-component are the probabilities that the process goes extinct in a fixed $A \subseteq X$, when starting from a vertex $x \in X$. The set of extinction probability vectors (obtained letting A vary among all subsets of X) is a subset of the set of the fixed points of the generating function of the branching random walk. In many cases there are only two possible extinction probability vectors and so far, in more complicated examples, only a finite number of distinct extinction probability vectors had been explicitly found. Whether a branching random walk could have an infinite number of distinct extinction probability vectors was an open question. We construct examples of branching random walks with uncountably many distinct extinction probability vectors.

Extinction in lower Hessenberg branching processes with countably many types

P. Braunsteins, The University of Queensland, Australia, p.braunsteins@uq.edu.au
S. Hautphenne, The University of Melbourne, Australia, sophiemh@unimelb.edu.au

We consider a class of branching processes with countably many types which we refer to as Lower Hessenberg branching processes. These are multitype Galton-Watson processes with typeset $\mathcal{X} = \{0, 1, 2, \ldots\}$, in which individuals of type i may give birth to offspring of type $j \leq i + 1$ only. For this class of processes, we study the set S of fixed points of the progeny generating function. In particular, we highlight the existence of a continuum of fixed points whose minimum is the global extinction probability vector $\boldsymbol{q}$ and whose maximum is the partial extinction probability vector $\tilde{\boldsymbol{q}}$. In the case where $\tilde{\boldsymbol{q}} = \mathbf{1}$, we derive a global extinction criterion which holds under second moment conditions, and when $\tilde{\boldsymbol{q}} < \mathbf{1}$ we develop necessary and sufficient conditions for $\boldsymbol{q} = \tilde{\boldsymbol{q}}$.

The probabilities of extinction in a branching random walk on a strip

P. Braunsteins, The University of Queensland, Australia, p.braunsteins@uq.edu.au
S. Hautphenne, The University of Melbourne, Australia, sophiemh@unimelb.edu.au

We consider a class of multitype Galton-Watson branching processes with countably infinite type set $\mathcal{X}$ whose mean progeny matrices have a block lower Hessenberg form. We study the probability of extinction in sets of types $A \subseteq \mathcal{X}$, $\boldsymbol{q}(A)$. In particular, we develop conditions for $\boldsymbol{q}(A)$ to be different from the global and partial extinction probability vectors. We present an iterative method to compute the vectors $\boldsymbol{q}(A)$, and investigate their location in the set of fixed points of the progeny generating vector. Finally, we extend some of these results to branching random walks on a tree.

Session 7.5 in Room M1 – Information and Learning in Stochastic Systems
Chair: Yuan Zhong

Experimenting in Equilibrium

S. Wager Graduate School of Business, Stanford University, swager@stanford.edu
K. Xu Graduate School of Business, Stanford University, kuangxu@stanford.edu

We study experimental design in large-scale stochastic systems with substantial uncertainty and structured cross-unit interference. We consider the problem of a platform that seeks to optimize supply-side payments p in a centralized marketplace where different suppliers interact via their effects on the overall supply-demand equilibrium, and propose a class of local experimentation schemes that can be used to optimize these payments without perturbing the overall market equilibrium. We show that, as the system size grows, our scheme can estimate the gradient of the platform's utility with respect to p while perturbing the overall market equilibrium by only a vanishingly small amount. We can then use these gradient estimates to optimize p via any stochastic first-order optimization method. These results stem from the insight that, while the system involves a large number of interacting units, any interference can only be channeled through a small number of key statistics, and this structure allows us to accurately predict feedback effects that arise from global system changes using only information collected while remaining in equilibrium.

Dynamic and stochastic knapsack problems: heuristics, asymptotic performance, and limit theorems

A. Arlotto, Duke University, United States, alessandro.arlotto@duke.edu
X. Xie, Duke University, United States, xinchang.xie@duke.edu

We study a dynamic and stochastic knapsack problem in which a decision maker is sequentially presented with n items with equal rewards and independent weights that are drawn from a known continuous distribution F. The decision maker seeks to maximize the expected reward she collects by including items in the knapsack while satisfying a capacity constraint, and while making terminal decisions as soon as each item weight is revealed. We consider different (near-) optimal policies. We study the total reward generated by each of these policies and we compare them in terms of their asymptotic mean, their asymptotic variance, and their limiting distributions. We observe that different asymptotically optimal policy may have different higher moments, and characterize policies that share the same limit theorems as the optimal one.

Deep Exploration via Randomized Value Functions

I. Osband, Deepmind, ian.osband@gmail.com
B. Van Roy, Stanford University, bvr@stanford.edu
D. Russo, Columbia, University, djr2174@gsb.columbia.edu
Z. Wen, Adobe Research, zwen@adobe.com

This talk will discuss some recent progress toward reliably efficient exploration in reinforcement learning systems. I'll introduce a new approach that generates sophisticated exploration by randomly perturbing value function estimates. This approach can be combined with common reinforcement learning algorithms, such as such as least-squares value iteration and temporal difference learning, which do not maintain a full model of the environment and instead aim to learn a parameterized representation of the value function. For tabular MDPs, I'll touch on a novel Bayesian analysis that proceed through stochastic dominance relations as well as some ongoing work that establishes worst-case regret bounds.

Individual Fairness in Hindsight

S. Gupta, Georgia Tech, USA, swatig@gatech.edu
V. Kamble, University of Illinois at Chicago, USA, kamble@uic.edu

Since many critical decisions impacting human lives are increasingly being made by algorithms, it is important to ensure that the treatment of individuals under such algorithms is demonstrably fair under reasonable notions of fairness. One compelling notion proposed in the literature is that of individual fairness (IF), which advocates that similar individuals should be treated similarly (Dwork et al. (2012)). Originally proposed for offline decisions, this notion does not, however, account for temporal considerations relevant for online decision-making. In this paper, we extend the notion of IF to account for the *time* at which a decision is made, in settings where there exists a notion of conduciveness of decisions as perceived by the affected individuals. We introduce two definitions: (i) fairness-across-time (FT) and (ii) fairness-in-hindsight (FH). FT is the simplest temporal extension of IF where treatment of individuals is required to be individually fair relative to the past as well as future, while in FH, we require a one-sided notion of individual fairness that is defined relative to *only* the past decisions. We show that these two definitions can have drastically different implications in the setting where the principal needs to learn the utility model. Linear regret relative to optimal individually fair decisions is inevitable under FT for non-trivial examples. On the other hand, we design a new algorithm: *Cautious Fair Exploration* (CaFE), which satisfies FH and achieves sub-linear regret guarantees for a broad range of settings. We characterize lower bounds showing that these guarantees are order-optimal in the worst case. FH can thus be embedded as a primary safeguard against unfair discrimination in algorithmic deployments, without hindering the ability to take good decisions in the long-run.

Session 7.6 in Room M2 – Statistics and Stochastic Processes
Chair: Sarat Moka

Modelling structural breaks in Autoregressive Time Series using parametric spectral discrimination tests

A. Grant, University of Cambridge, Cambridge, United Kingdom, andrew.grant@mrc-bsu.cam.ac.uk
L. Ma, Macquarie University, Sydney, Australia, lijing.ma@mq.edu.au
G. Sofronov, Macquarie University, Sydney, Australia, georgy.sofronov@mq.edu.au
J. R. Wishart, Macquarie University, Sydney, Australia, justin.wishart@mq.edu.au

It is quite common to have a time series whose structure changes abruptly. Identifying these change points and describing the model structure in the segments between these change points is of interest. In this talk, time series data is modelled assuming each segment is an Autoregressive time series with possibly different autoregressive parameters. This is achieved using two main steps. The first step is to use a cross entropy based estimation technique to identify these potential change points to segment the time series. Once these potential change points are identified, modified parametric spectral discrimination tests are used to validate the proposed segments. A numerical study is conducted to demonstrate the performance of the proposed method across various scenarios and compared against other contemporary techniques.

Algorithms for inferring parameters and hidden states of Markovian-regime-switching models with independent regimes for electricity prices

A. Lewis, The University of Adelaide, Australia, angus.lewis@adelaide.edu.au
N. Bean, The University of Adelaide, Australia, nigel.bean@adelaide.edu.au
G. Nguyen, The University of Adelaide, Australia, giang.nguyen@adelaide.edu.au

A popular model for wholesale electricity prices is the Markovian-regime switching (MRS) autoregressive process of order 1 (AR(1) process). The premise for using MRS models for economic time-series is that there may exist multiple regimes underlying the observed time-series, and depending on which regime the system is in, different characteristics are displayed. Due to the nature of electricity prices, specifically price spikes, it has become popular to use MRS models with *independent regimes*, which can capture spike behaviour with fewer regimes than their *dependent-regime* counterparts. We can think of these independent-regime MRS models for electricity prices as a collection of independent AR(1) processes, where at each time point only one of these processes is observed, and which process is observed is determined by a (hidden) Markov chain. We have developed new, computationally feasible, algorithms for the class of independent-regime MRS models to infer hidden states and find maximum-likelihood parameter estimates, given data. Our algorithms extend the forward-backward algorithm for hidden Markov models, which allows us to evaluate the likelihood and implement the EM algorithm to find maximum likelihood parameter estimates.

Distance covariance for discretized stochastic processes

H. Dehling, Ruhr-Universität Bochum, Germany, herold.dehling@rub.de
M. Matsui, Nanzan University, Japan, mmuneya@gmail.com
T. Mikosch, University of Copenhagen, Denmark, mikosch@math.ku.dk
G. Samorodnitsky, Cornell University, U.S.A., gs18@cornell.edu
L. Tafakori, RMIT University, Australia, laleh.tafakori@unimelb.edu.au

Given an iid sequence of pairs of stochastic processes on the unit interval we construct a measure of independence for the components of the pairs. We define distance covariance and distance correlation based on approximations of the component processes at finitely many discretization points. Assuming that the mesh of the discretization converges to zero as a suitable function of the sample size, we show that the sample distance covariance and correlation converge to limits which are zero if and only if the component processes are independent. To construct a test for independence of the discretized component processes we show consistency of the bootstrap for the corresponding sample distance covariance/correlation.

Unbiased Estimation of Reciprocal of the Mean for non-negative Random Variable

S. Moka, The University of Queensland, Australia, s.babumoka@uq.edu.au
D. Kroese, The University of Queensland, Australia, kroese@maths.uq.edu.au
S. Juneja, Tata Institute of Fundamental Research, India, juneja@tifr.res.in

We consider the problem of unbiased estimation of $1/\mathbb{E}[Z]$ for a non-negative random variable Z. In particular, we consider a Taylor expansion based unbiased estimator proposed by Blanchet et al. (2015) (which is in fact based on the seminal work Rhee and Glynn (2015)). This unbiased estimator – denote it by $V(x)$ – uses a Taylor expansion about a parameter x and truncates the infinite sum using an auxiliary random variable N. Then, $1/\mathbb{E}[Z]$ is estimated using the sample mean of K_x $i.i.d.$ samples of $V(x)$, where K_x is the maximum number of samples of $V(x)$ that can be generated within a given computing budget. In the literature, the *optimal* distribution of N for each x is well understood, and the optimal choice of x, however, is not known. In this paper, we focus on finding the optimal x for a given budget. As a consequence of our results, we show that it is optimal to choose $x = x^*$ where x^* is such that the average cost to generate one sample of $V(x^*)$ is equal to the entire budget. In other words, it is optimal to generate just one sample of $V(x^*)$ than generating K_x samples of $V(x)$ for any $x \neq x^*$. Furthermore, to establish valid confidence intervals for $V(x^*)$, we study its asymptotic distribution as the budget goes to infinity. We illustrate the applications of our findings using Gibbs hard-spheres process.

Session 7.7 in Rooms M5 and M6 – Matching and Compensation
Chair: Jussi Keppo

Efficient adjustment dynamics

S. Loertscher, University of Melbourne, Australia, simonl@unimelb.edu.au
E. Muir, Stanford University, USA, evmuir@stanford.edu

Assuming agents are endowed with one unit of a good and a commonly known maximum demand, arrive one after another and are privately informed about their values, we derive the efficient allocation and construct a mechanism that implements it. This setup permits definitions of buyers' and sellers' markets, breadth, depth, thickness, and liquidity of a market that relate only to efficiency and are thus mechanism-independent. We show that efficiency involves phases of market imbalance during which sellers or buyers are on the long side. The efficient adjustment dynamics to shocks, defined as out-of-equilibrium states, involve "fire trades" – the immediate liquidation of trades whose level exceeds the efficient storage threshold – followed by slow adjustments to a neutral state. Our framework also allows us to analyze optimal policy interventions, whereby a benevolent policymaker, observing the state of the order book, can improve next period's distribution at some cost to society.

Double Dipping of Two-Sided Platform Economy

S. Bhat, Institute of Operations Research and Analytics, National University of Singapore, satya.bhat@nus.edu.sg
J. Keppo, Business School, Institute of Operations Research and Analytics, and Institute of Data Science, National University of Singapore, keppo@nus.edu.sg
C. P. Teo, Business School and Institute of Operations Research and Analytics, National University of Singapore, bizteocp@nus.edu.sg

The role of platforms in breaking traditional monopolies, creating new markets, and resulting increase in social welfare is overwhelmingly positive. We examine a two-sided platform economy against a more stricter benchmark – an open platform that discloses its proprietary information. We show that, due to the big proprietary data the platforms control, they can be detrimental to both sides of the market resulting in social welfare losses compared to the benign open platform.

Approximating the Gittins Index in a Bayesian bandit

A. Lim, National University of Singapore, Singapore, andrewlim@nus.edu.sg

While the optimal policy for the Gittins–Whittle formulation of the multi-armed bandit is fully characterized in terms of the Gittins index, the Gittins index is notoriously difficult to compute for high dimensional Bayesian bandits. We develop a method for approximating the dynamics of the posterior using the Bayesian Central Limit Theorem and show how it can be used to approximate the Gittins index for high-dimensional Bayesian bandits. Comparisons of the Gittins–Whittle framework to Thompson sampling and the Upper Confidence Bound approach will be discussed, and applications of our approximation to Bayesian bandits where the rewards are mixtures will also be presented.

Mechanisms for Incentive-Compatible Information-Sharing in Informal Supply Chains

J. de Zegher, MIT, USA, jfz@mit.edu
I. Lo, Stanford University, USA, ilo@stanford.edu

Over 60%of global employment is informal, meaning that the majority of today's economy operates without formal contracts. In such informal working environments, people make decisions based on information that is exchanged through informal networks. This limits the information available to each member of the network and, therefore, their ability to make optimal decisions. In informal supply chains, members of the network also compete with each other, which disincentivizes the sharing of potentially valuable information. In this paper, we study the design of incentive-compatible information-sharing mechanisms for such informal supply chains. We incentivize information-sharing by selecting information disclosure policies that are contingent on the amount of information shared by each individual. We characterize optimal information disclosure policies and show that information sharing emerges as a subgame-perfect Nash equilibrium that is welfare-improving for all participants. Our model is informed by an information-sharing platform we deployed in informal palm oil supply chains in rural Indonesia, and we are currently designing a field experiment to test our proposed policies.

Session 7.8 in Rooms M7 and M8 – Ecology/Block Chain
Chair: Rhys Bowden

Applications of Bayesian inference, model selection, and experimental design in cell biology

D. J. Warne, Queensland University of Technology, Brisbane, Australia, david.warne@qut.edu.au
R. E. Baker, University of Oxford, Oxford, UK, ruth.baker@maths.ox.ac.uk
M. J. Simpson, Queensland University of Technology, Brisbane, Australia, matthew.simpson@qut.edu.au

Spatial models of collective cell behaviour are often based on reaction-diffusion models that describe a population of motile cells that proliferate. Various flux and source terms within the reaction-diffusion framework can be used to model different mechanisms. For example, motility may involve random motion (diffusion), adhesion, haptotaxis, chemokinesis and chemotaxis. In many applications, such as wound healing, it is not always clear which motility mechanisms are most important. Furthermore, the selection of appropriate growth functions is challenging since typical timescales of experimental protocols, based on cell culture assays, are too short to capture effects of contact inhibition of proliferation. We discuss how Bayesian approaches can provide significant insight into parameter inference, model selection, and experimental design for continuum models of cell motility and proliferation. We demonstrate our approach using microscopy data of PC-3 prostate cancer cells.

Stochastic Models of Multilevel Darwinian Populations

M. C. Nitschke, University of Adelaide, Australia, matthew.nitschke@adelaide.edu.au
A. Black, University of Adelaide, Australia, andrew.black@adelaide.edu.au
P. Rainey, Max Planck Institute for Evolutionary Biology, Germany, rainey@evolbio.mpg.de

The evolution of multicellular entities from unicellular ancestors signalled the emergence of a new level of biological organization. It is one of the major transitions in evolution. Despite its importance, there are many open questions surrounding the start of this transition. We investigate how this transition could have occurred by constructing stochastic models. These models describe an inherently stochastic process called Ecological Scaffolding, by which specific ecological conditions can cause individual cells to form groups and endow them with Darwinian properties; this means the groups show variation in character, differences in reproductive output, and heritability. This is implemented through a multilevel stochastic process which takes place over multiple time scales and includes a bottleneck and dispersal mechanism.

Polyp fiction: a stochastic fluid flow model for coral-algal symbiosis on the Great Barrier Reef

M. Wurm, The University of Adelaide, Australia, max.wurm@adelaide.edu.au
N. Bean, The University of Adelaide, Australia, nigel.bean@adelaide.edu.au
G. Nguyen, The University of Adelaide, Australia, giang.nguyen@adelaide.edu.au

Stochastic fluid models are highly versatile models which have seen use in a variety of different fields. In this context, we consider a fluid model where the level represents the algal density of coral on the Great Barrier Reef. This density of symbiotic algae is directly related to the threshold at which coral bleaching occurs, and is therefore of particular interest in the wake of global warming, a key contributor to coral bleaching on the Great Barrier Reef and all around the world. The survivability of the coral host can be approximated by a performance measure known as the time to mortality, the distribution of which can be obtained through the use of Laplace-Stieltjes transforms. In this talk, we present the derivation of this key performance measure, and how its distribution can be obtained using numerical inverse Laplace transform algorithms. Alongside the standard Abate-Whitt Euler method, we will look at a new method of numerical Laplace transform inversion, the concentrated matrix exponential (CME) method introduced in Horváth et al. (2019), and the performance of these methods in the context of the time to mortality.

Scaling and consensus in cryptocurrency blockchains

R. A. Bowden, University of Melbourne, Australia, rhys.bowden@unimelb.edu.au
A. E. Krzesinski, Stellenbosch University, South Africa, aek1@cs.sun.ac.za
P. G. Taylor, University of Melbourne, Australia, taylorpg@unimelb.edu.au

Bitcoin and other cryptocurrencies maintain a distributed global ledger in the form of a blockchain. Blocks are like pages in the ledger; each block is linked to the most recent prior block in the chain. Blocks are generated at random times at random mining computers in the Bitcoin network, then propagated to the other miners and the remainder of the network. Because blocks have a maximum size and are generated at a fixed average rate, there is a maximum rate at which transactions can be made, and this rate is orders of magnitude lower than other methods of electronic payment. In order to increase transaction throughput, some proposals suggest reducing the ratio of the average time to mine a block relative to the time it takes to propagate through the network. This results in different versions of the blockchain being present at different miners. We show different conditions under which these blockchains don't diverge, but stay close in some sense.

Friday 13:00–14:30

Session 8.1 in Rooms P1 and P2 – Simulation Optimization
Chair: Shane Henderson

Simple Bayesian Algorithms for Best Arm Identification

D. Russo, Columbia, University, djr2174@gsb.columbia.edu

This talk considers the optimal adaptive allocation of measurement effort for identifying the best among a finite set of options or designs. An experimenter sequentially chooses designs to measure and observes noisy signals of their quality with the goal of confidently identifying the best design after a small number of measurements. I will propose several simple Bayesian algorithms for allocating measurement effort, and by characterizing fundamental asymptotic limits on the performance of any algorithm, formalize a sense in which these seemingly naive algorithms are the best possible. Time permitting, I will also either present numerical experiments exhibiting performance surpassing competing approaches or discuss how some of these ideas may be extended to more complex problems.

Ranking and Selection under Input Uncertainty

D. Wu, Georgia Institute of Technology, dwu80@gatech.edu
E. Zhou, Georgia Institute of Technology, enlu.zhou@isye.gatech.edu

In stochastic simulation, input uncertainty (IU) is caused by the error in estimating the input distributions using finite real-world data. When it comes to simulation-based Ranking and Selection (R&S), ignoring IU could lead to the failure of many existing selection procedures. We study R&S under IU by allowing the possibility of acquiring additional data. Two classical R&S formulations are extended to account for IU: (i) for fixed confidence, we consider when data arrive sequentially so that IU can be reduced over time; (ii) for fixed budget, a joint budget is assumed to be available for both collecting input data and running simulations. New procedures are proposed for each formulation using the frameworks of Sequential Elimination and Optimal Computing Budget Allocation, with theoretical guarantees provided accordingly. Numerical results demonstrate the effectiveness of our procedures through a multi-stage production-inventory problem.

Knockout-Tournament Procedures for Large-Scale Ranking and Selection in Parallel Computing Environments

Y. Zhong, City University of Hong Kong, Hong Kong, yzhong4-c@my.cityu.edu.hk
L. Hong, Fudan University, China, hong_liu@fudan.edu.cn

On one hand, large-scale ranking and selection (R&S) problems require a large amount of computation. On the other hand, parallel computing environments that provide a large capacity for computation are becoming prevalent today and they are accessible by ordinary users. Therefore, solving large-scale R&S problems in parallel computing environments has emerged as an important research topic in recent years. However, directly implementing traditional stage-wise procedures and fully-sequential procedures in parallel computing environments may encounter problems, because either the procedures require too many simulation observations or the procedures' selection structures induce too many comparisons and too frequent communications among the processors. In this paper, inspired by the knockout-tournament arrangement of tennis Grand Slam tournaments, we develop new R&S procedures to solve large-scale problems in parallel computing environments. We show that, no matter whether the variances of the alternatives are known or not, our procedures can theoretically achieve the lowest growth rate on the expected total sample size with respect to the number of alternatives and thus are optimal in rate. Moreover, common random numbers (CRNs) can be easily adopted in our procedures to further reduce the total sample size. Meanwhile, the comparison time in our procedures is negligible compared to the simulation time, and our procedures barely request for communications among the processors.

Fixed-confidence, fixed-tolerance guarantees for ranking-and-selection procedures

D. J. Eckman, Cornell University, USA, dje88@cornell.edu
S. G. Henderson, Cornell University, USA, sgh9@cornell.edu

Ranking and selection (R&S) procedures are used to select, from a finite number of simulated systems, the one with the largest performance measure. With only a small number of exceptions, R&S procedures provide a frequentist *probability of correct selection* (PCS) guarantee under an indifference-zone assumption whereby the best system is assumed to be strictly better than the second-best system by at least a known positive amount. R&S procedures that make this assumption typically provide no guarantee when the assumption fails, though empirically they seem to perform well. We argue for the stronger *probably approximately correct* (PAC) guarantee in use in, e.g., the related area of exploratory bandits. We provide sufficient conditions that "lift" the weaker PCS guarantee to the stronger PAC guarantee, and discuss Bayesian formulations.

Session 8.2 in Room P3 – Resource Allocation Models
Chair: Alexander Stolyar

Stability and moment bounds under utility-maximising service allocations

S. Shneer, Heriot-Watt University, UK, V.Shneer@hw.ac.uk
A. Stolyar, University of Illinois at Urbana-Champaign, USA, stolyar@illinois.edu

We study networks of interacting queues governed by utility-maximising service-rate allocations in both discrete and continuous time. We establish stability and some steady-state moment bounds under natural conditions and rather weak assumptions on utility functions. These results are obtained using direct applications of Lyapunov-Foster- type criteria, and apply to a wide class of systems, including those for which fluid limit-based approaches are not applicable.

Improved queue-size scaling for input-queued switches via graph factorization

J. Xu, Fuqua School of Business, Duke University, USA, jx77@duke.edu
Y. Zhong, Booth School of Business, University of Chicago, USA, yuan.zhong@chicagobooth.edu

We consider the scaling of the expected total queue size in an $n \times n$ input-queued switch, as a function of both the load ρ and the system scale n. In the regime $1-\rho = O(1/n)$, we provide a new class of scheduling policies under which the expected total queue size scales as $O(n(1-\rho)^{-4/3}\log(1-\rho)^{-1})$. This improves over the previously best-known scaling $O(n^{1.5}(1-\rho)^{-1}\log(1-\rho)^{-1}$ when $\Omega(n^{-1.5}) \leq 1-\rho \leq O(n^{-1})$. A key ingredient in our method is a tight characterization of the largest k-factor of a random bipartite multigraph, which may be of independent interest.

Information and Memory in Dynamic Resource Allocation

K. Xu Graduate School of Business, Stanford University, kuangxu@stanford.edu
Y. Zhong Booth School of Business, University of Chicago, Yuan.Zhong@chicagobooth.edu

We propose a general framework, dubbed *Stochastic Processing under Imperfect Information (SPII)*, to study the impact of information constraints and memories on dynamic resource allocation. The framework involves a Stochastic Processing Network (SPN) scheduling problem in which the decision maker may access the system state only through a noisy *channel*, and resource allocation decisions must be carried out through the interaction between an encoding policy (who observes the state) and allocation policy (who chooses the allocation) situated on the two ends of the channel. Applications in the management of large-scale data centers and human-in-the-loop service systems are among our chief motivations.
We quantify the degree to which information constraints reduce the size of the *capacity region* in general SPNs, and how such reduction depends on the amount of *memories* available to the encoding and allocation policies. Using a novel metric, *capacity factor*, our main theorem characterizes the reduction in capacity region (under "optimal" policies) for all non-degenerate channels, and across almost all combinations of memory sizes. Notably, the theorem demonstrates, in substantial generality, that (1) the presence of a noisy channel always reduces capacity, (2) more memories for the allocation policy always improve capacity, and (3) more memories for the encoding policy have little to no effect on capacity. Finally, all of our positive (achievability) results are established through constructive, implementable policies.
Our proof program involves the development of a host of new techniques, largely from first principles, by combining ideas from information theory, learning and queueing theory. As a sub-module of one of the policies proposed, we create a simple yet powerful generalization of the Max-Weight policy, in which individual Markov chains are selected dynamically, in a manner analogous to how schedules are used in a conventional Max-Weight policy.

Discrete-time TASEP with holdback

S. Shneer, Heriot-Watt University, UK, V.Shneer@hw.ac.uk
A. Stolyar, University of Illinois at Urbana-Champaign, USA, stolyar@illinois.edu

We study a discrete-time interacting particle process, which can be called a totally asymmetric simple exclusion process with holdback (TASEP-H). The “holdback” refers to the property that the probability of a particle jumping forward to a vacant site depends on the presence of a particle immediately “behind” it. This is a basic model of a communication network with packets moving along a sequence of nodes, under a “standard” random access algorithm. We obtain results on “typical” invariant distributions of the process.

Session 8.3 in Room P4 – Statistical Inference in Queueing Systems
Chair: Liron Ravner

A Survey of Parameter and State Estimation in Queues

A. Asanjarani, The University of Auckland, New Zealand, azam.asanjarani@auckland.ac.nz
Y. Nazarathy, The University of Queensland, Australia, y.nazarathy@uq.edu.au
P. Taylor, The University of Melbourne, Australia, taylorpg@unimelb.edu.au

We present a broad literature survey of parameter and state estimation for queueing systems. Our approach is activity based. We categorise results based on the modelling activity that can be done. Categories include: prediction, probing, modelling of non-interacting systems (infinite servers), queue inference engine approaches, likelihood estimation approaches, Bayesian approaches, solving inverse problems, handling discrete sampling, using queueing fundamentals such $L = \lambda W$, information mining, and control, design, and uncertainty quantification.

Statistical estimation of the input to a queue by Poisson probing

L. Ravner, University of Amsterdam, The Netherlands, l.ravner@uva.nl
O. Boxma, Eindhoven University of Technology, The Netherlands, O.J.Boxma@tue.nl
M. Mandjes, University of Amsterdam, The Netherlands, m.r.h.mandjes@uva.nl

In many queueing scenarios the arrival process of work to the queue may be unknown and so statistical inference by observing system performance is called for. For example, a communication network with a bottleneck link that has multiple unknown sources of input. This is often a challenging task as time dependent samples from a queue have an intractable joint distribution. In this work we suggest a tractable approach that relies on sampling the workload process at random times according to an external Poisson process. We assume that the input is a Lévy process, which includes the special case of the classical M/G/1 model. We construct a method-of moments based estimator for the characteristic exponent (generating function) of the input distribution. The estimator relies on transient properties of the queue rather than steady-state approximations. Verifiable conditions for consistency and asymptotic normality are provided, along with explicit expressions for the asymptotic variance. We show that an estimator satisfying these conditions can be constructed for a class of models including the M/G/1 queue. We will further discuss how this method may be useful in more elaborate systems such as tandem and tree networks.

Hypothesis testing in queueing

M. Mandjes, University of Amsterdam, the Netherlands, m.r.h.mandjes@uva.nl
L. Ravner, University of Amsterdam, the Netherlands, l.ravner@uva.nl

In this talk I'll discuss a set of hypothesis testing problems in queueing. The leading example concerns a classical $M/G/1$ queue in which the workload is observed at Poisson epochs, where H_0 and H_1 correspond to specific combinations of arrival rates and service-time distributions. As the likelihood ratio cannot be computed in closed form, we consider two other test statistics: one is based on the probability of seeing an empty queue conditional on the workload at the previous inspection epoch, while the other uses ‘busy periods at Poisson epochs’. For both statistics we point out how to calibrate the thresholds for the corresponding test. This is joint work with Liron Ravner (University of Amsterdam).

Session 8.4 in Room P5 – Recent Developments and Advances in Markov Decision Processes
Chair: Sandjai Bhulai & Floske Spieksma

Asymptotically optimal control for Markov Decision Processes (MDP) under side constraints

A. Burnetas, National and Kapodistrian University of Athens, Greece, aburnetas@math.uoa.gr
O. Kanavetas, Rutgers University, USA, okanavetas@business.rutgers.edu
M. Katehakis, Rutgers University, USA, mnk@rutgers.edu

We consider the problem of adaptive control for Markov Decision Processes (MDP), under side constraints, when there is incomplete information for the transition probabilities and its rewards. Under suitable irreducibility assumptions for the MDP we construct adaptive policies that maximize the rate of convergence of realized rewards to that of the optimal (non adaptive) policy under complete information. We discuss applications and computational results for queuing control problems and contextual bandits.

Data-driven consumer debt collection via machine learning and approximate dynamic programming

R. van de Geer, Vrije Universiteit Amsterdam, The Netherlands, r.vande.geer@vu.nl
Q. Wang, Amsterdam Business School, The Netherlands, Q.Wang@uva.nl
S. Bhulai, Vrije Universiteit Amsterdam, The Netherlands, s.bhulai@vu.nl

We develop and test a framework for the data-driven scheduling of outbound calls made by debt collectors. These phone calls are used to persuade debtors to settle their debt, or to negotiate payment arrangements in case debtors are willing, but unable to repay. We determine on a daily basis which debtors should be called to maximize the amount of delinquent debt recovered in the long term, under the constraint that only a limited number of phone calls can be made each day. Our approach is to formulate a Markov decision process and, given its intractability, approximate the value function based on historical data through the use of state-of-the-art machine learning techniques. Specifically, we predict the likelihood with which a debtor in a particular state is going to settle its debt and use this as a proxy for the value function. Based on this value function approximation, we compute for each debtor the marginal value of making a call. This leads to a particularly straightforward optimization procedure, namely, we prioritize the debtors that have the highest marginal value per phone call. We validate our proposed methodology in a controlled field experiment conducted with real debtors. The results show that our optimized policy substantially outperforms the current scheduling policy that has been used in business practice for many years. Most importantly, our policy collects more debt in less time, whilst using substantially fewer resources — leading to a large increase in the amount of debt collected per phone call.

Optimal Policies for Stochastic Clearing Systems with Time-Dependent Delay Penalties

Q.-M. He, University of Waterloo, Canada, q7he@uwaterloo.ca
J. H. Bookbinder, University of Waterloo, Canada, jbookbinder@uwaterloo.ca
Q. Cai, University of Waterloo, Canada, saulcai@hotmail.com

We study stochastic clearing systems that are characterized by a discrete-time Markovian input process, and an output mechanism that intermittently and instantaneously clears the system partially or completely. The outstanding inputs in the system are recorded in a sequence, which allows us to keep track of the delays and quantities of each input separately. The decision to clear the system depends on both quantities and individual delays. Clearing the system incurs a fixed cost, and outstanding inputs are charged a delay penalty, which is a general increasing function of the quantities and delays of individual inputs. We model the system as a tree-structured Markov decision process over both finite and infinite horizon. We show properties of the optimal clearing policies (e.g., the on-off type or the threshold type), and develop algorithms to compute parameters of those optimal policies. We conduct a numerical analysis on the impact of the nonlinear cost function, the comparison of the optimal policy and the classical hybrid policy, and the impact of the state of the input process. Our experiments demonstrate that i) the classical linear approximation of the cost function can lead to significant performance differences; ii) the classical hybrid policy may perform poorly (as compared to the optimal policy); and iii) the consideration of the state of the input process makes significant improvement in system performance.

Bounds for threshold and switching curve optimal policies

F. M. Spieksma, Universiteit Leiden, The Netherlands, spieksma@math.leidenuniv.nl

Consider a single server queueing system with service control. Customers arrive according to a Poisson process and require an exponentially distributed amount of service. Control is exercised by selecting the speed at which the server operates: a low or a high one. Switching the server to the high speed gives high operating cost. On the other hand, serving at low cost tends to increase the holding cost for customers in queue.
It is known that the optimal control policy minimising the expected average cost per unit time is a threshold policy, and the value function is convex as a function of state, if e.g. the holding cost is convex as a function of state. This can be shown by applying value iteration to the uniformised time-discretised control problem.
It does not seem to be universally known though, that choosing appropriate initial functions in the value iteration step, yields upper and lower bounds for both the threshold optimal policy and the value function in each iteration step. Moreover, the bounds converge monotonically to optimal threshold and value function, so that the region where the optimal policy is not known is non-increasing in the iteration step. This allows more efficient numerical computation.
The method is not restricted to one-dimensional queueing models. As a two-dimensional extension, we also investigate the 2-competing queues model with quadratic holding cost per unit time, as a function of the number of customers in each of the queues. Numerical results indicate that the switching curve is linear, with slope equal to the proportion of the cost reduction per unit time due to serving one type versus serving the other type. As such, it has some similarities to the $c\mu$-rule in the case of linear cost. An interesting open problem is the proof of this interesting behaviour of the switching curve.

Session 8.5 in Room M1 – Learning, Optimization, and Applied Probability
Chair: Harsha Honnappa

Discriminative Learning via Adaptive Questioning

A. Bassamboo, Northwestern University, a-bassamboo@kellogg.northwestern.edu
S. Juneja, Tata Institute of Fundamental Research, Mumbai, juneja@tifr.res.in
A. Zeevi, Columbia University, assaf@gsb.columbia.edu

We consider the problem of designing an adaptive sequence of questions that optimally classify a candidate's ability into one of several categories or discriminative grades. A candidate's ability is modelled as an unknown parameter, which, together with the difficulty of the question asked, determines the likelihood with which s/he is able to answer a question correctly. The learning algorithm is only able to observe these noisy responses to its queries. We consider this problem from two main perspective. A fixed confidence-based IZ-PAC framework, that in our setting seeks to arrive at the correct ability discrimination at the fastest possible rate while guaranteeing that the probability of error is less than a pre-specified and small IZ. In this setting we develop lower bounds on any sequential questioning strategy and arrive at algorithms that match these lower bounds to the first order. A complementary formulation is one where the number of questions asked is fixed and the aim is to minimize the probability of incorrect ability discrimination. In this so-called fixed budget setup the problem of designing optimal (or asymptotically optimal) algorithms is notoriously difficult. We adopt a slight modification via an indifference zone that allows the derivation of lower bounds and the design of algorithms whose performance asymptotically achieve these bounds. Our results leverage optimization problems defined via the lower bound derivations and give rise to several structural insights on the format and diversity of difficulty levels in the optimal discriminative exam.

Distributed Learning in Multi-Agent Systems: Optimization and Reinforcement Learning

T. T. Doan, Georgia Institute of Technology, USA, thinh.doan@isye.gatech.edu
S. T. Maguluri, Georgia Institute of Technology, USA, siva.theja@gatech.edu
J. Romberg, Georgia Institute of Technology, USA, jrom@ece.gatech.edu

The rapid development of low-cost sensors, smart devices, communication networks, and learning algorithms has enabled data-driven decision making in large-scale multi-agent systems such as mobile robotic networks, smart grids, and autonomous systems. The key challenge in these systems is in handling the vast quantities of information shared between the agents in order to find an optimal or near-optimal policy that maximizes an objective function. This task needs to be done under computation and communication constraints. Distributed consensus-based gradient (DCG) method is a popular low cost approach that is amenable to real time implementation. In this talk, I will present some results characterizing the convergence rates of DCG algorithms, and how the optimal rates can be obtained even in presence of communication constraints such as limited bandwidth. I will also demonstrate how these results suggest a natural distributed TD algorithm for reinforcement learning, and characterize its convergence rates.

Stochastic Modeling and Decisions for Online-to-offline Supermarket Retails

J. Cao, University of California, Berkeley, jycao@berkeley.edu
Z. Shen, University of California, Berkeley, maxshen@berkeley.edu
Z. Zheng, University of California, Berkeley, zyzheng@berkeley.edu

Online-offline supermarket chains, with a hybrid of both traditional offline retailing and online ordering followed by fast delivery, are seeing rapid growth in the world and particularly in Chinese major cities. With the aid of data analytics and supply chain upgrades, such online-offline supermarkets are able to provide superb fresh food that are guaranteed to be produced on the same day of sales. This service, in turn, triggers high-volume high-volatility demands and incurs the need to further enhance intra-day operational decisions. In this talk, we discuss stochastic models and optimization for several such operational decisions, including intra-day replenishment strategies and dynamic replenishment delivery routes design across multiple stores.

Stochastic Trust Region Methods ASTRO and ASTRO-DF: Convergence, Complexity, and Numerical Performance

D. Vasquez, Purdue University, Department of Statistics, dvasque@purdue.edu
R. Pasupathy, Purdue University, Department of Statistics, pasupath@purdue.edu
S. Shashaani, Purdue University, Department of Industrial & Systems Engineering, sshasha2@ncsu.edu

ASTRO is an adaptive sampling trust region algorithm designed to solve (unconstrained) smooth stochastic optimization problems with a first-order oracle constructed using a Monte Carlo simulation or a large dataset of “scenarios”. Like its derivative-free counterpart ASTRO-DF, ASTRO adaptively constructs a stochastic local model, which is then optimized and updated iteratively. Crucially and unlike ASTRO-DF, however, the extent of sampling in ASTRO is directly proportional to the inverse squared norm of the sampled gradient at the incumbent point, leading to clearly observable gains in numerical performance. The sequence of (true) gradient norms measured at ASTRO’s iterates converges almost surely to zero. We also characterize a work complexity result that expresses ASTRO’s convergence rate in terms of the total number of oracle calls (as opposed to the total number of iterations). This appears to be one of the few results of its kind and confirms the strong numerical performance observed in practice.

Session 8.6 in Room M2 – Large-Scale Inference and High-Dimensional Statistics Chair: Wen Sun

High-dimensional regression with discrete optimization

H. Hazimeh, Massachusetts Institute of Technology, USA, hazimeh@mit.edu
R. Mazumder, Massachusetts Institute of Technology, USA, rahulmaz@mit.edu
P. Radchenko, The University of Sydney, Australia, peter.radchenko@sydney.edu.au

We will discuss recent applications of discrete optimization techniques in high-dimensional regression. In particular, we will review the recently proposed mixed integer optimization implementation of the discrete Dantzig Selector. The estimator minimizes the number of nonzero regression coefficients, subject to a budget on the maximal absolute correlation between the features and residuals. It can be expressed as a solution to a mixed integer linear optimization problem, a computationally tractable framework that delivers provably optimal global solutions. We will also discuss applications of mixed integer optimization in high-dimensional linear regression with group structure, as well as high-dimensional additive modelling. In addition, we will consider a regularized version of the best subset selector, and investigate its advantages in the low signal regimes.

Community Detection via L-1 Fused Penalty

Y. Choi, National University of Singapore, Singapore, stachoiy@nus.edu.sg
V. Tan, National University of Singapore, Singapore, vtan@nus.edu.sg
Z. Liu, National University of Singapore, Singapore, a0109983@u.nus.edu

In recent years, community detection has been an active research area in various fields including machine learning and statistics. While a plethora of works has been published over the past few years, most of the existing methods depend on a predetermined number of communities. Given the situation, determining the proper number of communities is directly related to the performance of these methods. Currently, there does not exist a golden rule for choosing the ideal number, and people usually rely on their background knowledge of the domain to make their choices. To address this issue, we propose a community detection method that is equipped with data-adaptive methods of finding the number of the underlying communities. Central to our method is fused l-1 penalty applied on an induced graph from the given data. The proposed method shows promising results.

Privacy Preserving Integrative Regression Analysis of High-dimensional Heterogeneous Data

Y. Xia, Fudan University, Shanghai, xiayin@fudan.edu.cn

Meta-analyzing multiple studies, enabling more precise estimation and investigation of generalizability, is important for evidence based decision making. Integrative analysis of multiple heterogeneous studies is, however, highly challenging in the high dimensional setting. The challenge is even more pronounced when the individual level data cannot be shared across studies due to privacy concerns. Under ultra high dimensional sparse regression models, we propose in this talk a novel integrative estimation procedure by aggregating and debiasing local estimators (ADeLE), which allows us to base solely on the derived data to perform estimation with general loss functions. The ADeLE procedure accommodates between study heterogeneity in both the covariate distribution and model parameters, and attains consistent variable selection. Furthermore, the prediction and estimation errors incurred by aggregating derived data is negligible compared to the statistical minimax rate. In addition, the ADeLE estimator is shown to be asymptotically equivalent in prediction and estimation to the ideal estimator obtained by sharing all data. The finite-sample performance of the ADeLE procedure is studied via extensive simulations. We further illustrate the utility of the ADeLE procedure to derive phenotyping algorithms for coronary artery disease using electronic health records data from multiple disease cohorts.

Estimating the probability of emergence of a novel disease strain against an endemic background

R. Cope, The University of Adelaide, Australia, robert.cope@adelaide.edu.au
M. Meehan, James Cook University, Australia, michael.meehan1@jcu.edu.au

Branching process approximations are often used to analyse the early dynamics of a disease introduced into an infection-naive population. The theory is well developed, including on networks, or when infectious individuals have (random) heterogeneous reproductive capacities. We adapt this theory to analyse the emergence of a novel disease strain competing against an existing endemic disease, when the two strains share cross-immunity. This allows us to estimate the probability the new strain will become established, based on the relative fitness of the resident and invading strains. We evaluate the impact of different control measures on the probability of emergence of the new strain, and find that certain control regimes can facilitate its emergence, relative to the control-free case. We also present a simulation study demonstrating the impact of different parameters on the accuracy of the branching process approximation relative to exact simulation of the full disease process. This work will enhance our understanding of the emergence of new disease strains, and particularly antimicrobial drug resistant strains, which present one of the most significant challenges to public health globally.

Session 8.7 in Rooms M5 and M6 – Storage Processes and Markov Chains
Chair: Jaron Sanders

Steady-state optimization of an exhaustive Lévy storage process with intermittent output and random output rate

R. Jacobovic, The Hebrew University of Jerusalem, Israel, royi.jacobovic@mail.huji.ac.il
O. Kella, The Hebrew University of Jerusalem, Israel, Offer.Kella@huji.ac.il

Consider a regenerative storage process with a nondecreasing Lévy input (subordinator) such that every cycle may be split into two periods. In the first (vacation) the output is shut off and the workload is accumulated. This continues until some stopping time. In the second, the process evolves like a subordinator minus a positive drift (output rate) until it hits the origin. In addition, we assume that the output rate of every busy period is a random variable which is determined at the beginning of this period. For example, at each period, the output rate may depend on the workload level at the beginning of the corresponding busy period. We derive the Laplace–Stieltjes transform of the steady state distribution of the workload process and then apply this result to solve a steady-state cost minimization problem with holding, setup and output capacity costs. It is shown that the optimal output rate can be determined as a nondecreasing deterministic function of the workload level at the beginning of the corresponding busy period. This function is given explicitly up to a solution of a certain equation. Finally, numerical results are provided for special cases.

Sojourn time distribution in fluid queues

E. Deiana, Université de Namur, Belgium, eleonora.deiana@unamur.be
G. Latouche, Université libre de Bruxelles, Belgium, latouche@ulb.ac.be
M.-A. Remiche, Université de Namur, Belgium, marie-ange.remiche@unamur.be

A fluid flow $\{X(t), S(t)\}_{t\geq 0}$ represents the behavior of a reservoir content which evolves linearly following some input and output rates modulated by a background Markov process $S(t)$. Fluid flows have applications in telecommunication and computer systems, where they model the amount of data entering and being processed by a server. When studying these systems the attention is generally focused on the stationary distribution of the level, while less attention is paid to the distribution of the sojourn time.
In a traditional queueing system, where clients arrive and are served individually, the sojourn time of a typical client is the time spent inside the system, from the instant of its arrival until its departure. In a fluid flow, the clients have infinitesimal size and they arrive and leave the system continuously. The definition of the sojourn time is more involved, and a different approach is required. We proceed in two steps. We first compute the stationary distribution of the buffer at arrival instants. This is based on the definition of a new clock $\phi(y)$, which is the epoch when a total amount of y units of fluid entered the buffer. Secondly, we compute the transform of the time spent to empty the buffer.
The *new clock* approach is readily extended to systems of increased complexity. For example, in the case of a finite buffer, we define the new clock taking into account the amount of fluid which is lost because of buffer overflow.

Limit Theorems for Markov Chains on Sparse Graphs

A. Ganguly, Brown University, U.S., ankan_ganguly@brown.edu
K. Ramanan, Brown University, U.S., kavita_ramanan@brown.edu

The mean-field approximation is a powerful tool in the analysis of interacting particle systems on large graphs. However, although the mean-field approximation is asymptotically exact on large dense graphs, it is inaccurate on sparse graphs. In the first half of this talk, I will discuss several motivating examples of interacting particles on large sparse graphs for which the dynamics of individual particles are of interest. I will also introduce a few alternatives to the mean field model. In the second half of the talk, I will introduce a novel characterization of the dynamics of a typical particle on certain sparse graphs. This characterization relies on a Gibbs-like conditional independence property that holds with great generality. This talk will be based on my work with Kavita Ramanan.

Clustering in Block Markov Chains

J. Sanders, Delft University of Technology, The Netherlands, j.sanders@tudelft.nl
A. Proutière, KTH Royal Institute of Technology, Sweden, alepro@kth.se
S.-Y. Yun, Korea Advanced Institute of Science and Technology, South Korea, yunseyoung@kaist.ac.kr

In this talk, we will consider cluster detection in Block Markov Chains (BMCs). These Markov chains are characterized by a block structure in their transition matrix. More precisely, the n possible states are divided into a finite number of K groups or clusters, such that states in the same cluster exhibit the same transition rates to other states. One observes a trajectory of the Markov chain, and the objective is to recover, from this observation only, the (initially unknown) clusters. For BMCs, we have devised a clustering procedure that accurately, efficiently, and provably detects the clusters. We first derive a fundamental information-theoretical lower bound on the detection error rate satisfied under any clustering algorithm. This bound identifies the parameters of the BMC, and trajectory lengths, for which it is possible to accurately detect the clusters. We next develop two clustering algorithms that can together accurately recover the cluster structure from the shortest possible trajectories, whenever the parameters allow detection. These algorithms thus reach the fundamental detectability limit, and are optimal in that sense.

Session 8.8 in Rooms M7 and M8 – Stochastic Modeling and Control in Healthcare and Service Operations
Chair: Pengyi Shi

Managing Time-varying Arrivals in Service Systems: A Dual Approach

X. Peng, Indiana University, Bloomington, IN, USA, xp1@iu.edu

This paper studies the optimal admission control problem when the arrival rate is itself a time-dependent stochastic process. In the baseline model, we consider the classical admission control problem for a single-server queue (without forecast updates). The system incurs a waiting cost of h per unit time for each customer waiting in the queue, whereas it incurs a one-time penalty of p if the customer is not admitted to the system. We allow the arrival process to be a general stochastic process. The system seeks to minimize the total expected cost incurred within a day. We formulate the problem as a finite-horizon stochastic control problem and derive its dual. The solution to the dual problem can be interpreted as the expected cost incurred by the incoming customer at the moment. We refer it to as the shadow price process. We show that the shadow price process has two components: A deterministic part and a martingale. The first component is a deterministic process that characterizes the impact of the time-dependence effect of the arrival process. In addition, we show that the deterministic component is the solution to a deterministic optimal control problem that minimizes the total cost in the fluid analog of the original system. The second component of the shadow price process is a martingale that characterizes the impact of the uncertainty in arrivals. The evolution of the martingale component characterizes how the shadow price process evolves as the (previously unknown) information is revealed due to the realization of the arrivals. We then extend the admission control problem to the setting where the system is allowed to update its arrival forecast. We consider two special scenarios. In the first scenario, the system forms a prior belief about the distribution of the arrival process. As the arrivals are realized during the day of operation, the system manager revises his belief about the arrival process using Bayesian updating continuously. In the second scenario, we consider a case when the system uses outside information to update the belief about the arrivals. In both settings, we show how the forecast update affects the admission decision through the evolution of the martingale. We then use the characterization of the martingale component to quantify the value of the arrival forecast.

To Pool or Not to Pool: Queueing Design for Large-Scale Service Systems with Customer Abandonment

P. Cao, University of Science and Technology of China, China, pcao@ustc.edu.cn
S. He, National University of Singapore, Singapore, heshuangchi@nus.edu.sg
J. Huang, Chinese University of Hong Kong, China, junfeih@cuhk.edu.hk
Y. Liu, North Carolina State University, yliu48@ncsu.edu

There are two basic queue structures commonly adopted in service systems: the pooled structure where waiting customers are organized into a single queue served by a group of servers and the dedicated structure where each server has her own queue. Although the pooled structure, known to minimize the servers' idle times, is widely used in large-scale service systems, this study reveals that the dedicated structure, along with the join-the-shortest-queue routing policy, could be more advantageous for improving some service levels, such as the probability of a customer's waiting time being within a delay target. The servers' additional idleness resulted from the dedicated structure will be negligible when the system has many servers. Using a fluid model substantiated by asymptotic analysis, we provide a performance comparison between the two structures for a moderately overloaded queueing system with customer abandonment. We intend to help service system designers answer the following questions: To achieve a specified service level, which queue structure will be more cost-effective? How many servers can be saved by converting one structure into the other? Aside from structure design, our results are also of practical value for performance analysis and staffing deployment.

Optimal Scheduling of Proactive Care with Customer Degradation

Y. Hu, Columbia University, USA, YHu22@gsb.columbia.edu
C. Chan, Columbia University, USA, cwchan@gsb.columbia.edu
J. Dong, Columbia University, USA, jing.dong@gsb.columbia.edu

Healthcare is a limited resource environment where the most severe patients are typically prioritized when capacity is scarce. However, there has been a growing interest in proactive (or preventative) care. On one hand, providing care for patients when they are less critical could mean that fewer resources are necessary to return them to a healthy, stable state. On the other hand, utilizing limited capacity for patients who may never need care in the future takes the capacity away from other more critical patients who need it now. To understand this tension, we propose a multi-server queueing model with two patient classes: moderate and urgent. A moderate patient who does not receive treatment may recover and leave or may deteriorate and become an urgent patient. In this setting, we characterize how moderate and urgent patients should be prioritized for care when proactive care for moderate patients is an option. The analysis replies on several innovative applications of optimal control theory.

Integrating New Diagnostic Tests Into Emergency Department Workflow

J. Helm, Indiana University, USA, helmj@indiana.edu
P. Shi, Purdue University, USA, shi178@purdue.edu
S. Heese, North Carolina State University, hsheese@ncsu.edu
A. Mitchell, IU School of Medicine, alimitch@iu.edu

In medical research, new diagnostic tests are developed and evaluated solely on their efficacy in detecting an illness. However, ignoring the workload impact of introducing new tests into existing workflow can create barriers to adoption, particularly in busy units such as Emergency Department (ED). In collaboration with an ED physician, we develop an analytical framework for evaluating the workload impact of adopting new tests to bridge the gap between medical research and clinical workflow. We derive useful structural properties on when and which patient to receive the new test, and develop a decomposition algorithm to solve the problem efficiently when the number of patient classes is large. As a showcase, we apply our framework to a new test, D-dimer, for diagnosing Pulmonary Embolism in EDs to understand what characteristics make adopting this new test feasible and how to best integrate it into the ED workflow.

Friday 15:00–16:30

Session 9.1 in Rooms P1 and P2 – Convex Optimization in Applied Probability
Chair: Chaithanya Bandi

Near Optimal Control of a Ride-Hailing Platform via Mirror Backpressure

P. Qian, Columbia University, United States, pq2124@columbia.edu
Y. Kanoria, Columbia University, United States, ykanoria@columbia.edu

Ride-hailing platforms need to match supply and demand so as to maximize the rate of payoff generation while respecting the geographical flow constraints, namely, that the rate at which vehicles arrive at a location must equal the rate at which they leave that location. The platform's control levers include: (i) *entry control*, i.e. it can choose not to serve some customers, (ii) *dynamic pricing*, i.e. it can set prices for each ride depending on its origin and destination, and (iii) *assignment rule*, i.e. it can choose from which neighboring location to dispatch a car for pickup.
We consider two settings depending on whether pricing is an available lever to the platform. In joint-entry-assignment (JEA) setting, the platform can only use levers (i) and (iii) (for example, Didi in China does not use dynamic pricing). In joint-pricing-assignment (JPA) setting, the platform can use levers (ii) and (iii) (e.g. most ride-hailing platforms in North America). We introduce a novel family of *Mirror Backpressure* (MBP) platform control policies which are simple, practical, and do not require prior knowledge of the demand arrival rates. Key challenges include that serving a customer reduces supply availability at the dispatch location but increases it at the dropoff location ("supply externalities"), and that the number of vehicles at any location can never be negative ("no-underflow constraints"). Mirror Backpressure generalizes the backpressure policy such that it executes mirror descent, allowing us to address these challenges. MBP loses at most an $O(K/T + 1/K)$ fraction of the achievable payoff in the JEA setting and $O(\sqrt{K/T + 1/K})$ fraction in the JPA setting, where K is the number of vehicles and T is the horizon. Simulation results in a realistic environment support our theoretical findings.

Dynamic Pricing in Two Sided Queues

S. Varma, Georgia Institute of Technology, USA, sushil@isye.gatech.edu
P. Bumpensanti, Georgia Institute of Technology, USA, pornpawee@gatech.edu
H. Wang, Georgia Institute of Technology, USA, he.wang@isye.gatech.edu
S. T. Maguluri, Georgia Institute of Technology, USA, siva.theja@gatech.edu

We study two sided queuing systems in which, unlike the traditional queuing systems, customers and server both arrive and wait in the queue until they are matched. As soon as a customer is matched with a server, they leave the system and the service is assumed to be complete instantaneously. The major issue with these systems is that a constant customer and server arrival rates cannot make the system positive recurrent. We deal with this issue by employing the idea of dynamic pricing. We formulate our system as an average cost Markov decision process to maximize the revenue earned by the system operator, where the operator can choose a pricing policy from a set of stationary pricing policies. We proved that under the optimal pricing policies, the customer prices and server prices are monotonic in the queue length. We also showed that if we approximate the value function by a polynomial in queue length, then the policy with optimal coefficients of the polynomial gives an upper bound on the revenue. Moreover, asymptotically, the system has a positive deterministic component and a negative stochastic component. A two price policy leads to $n^{1/3}$ scale effect of variability. Lastly, we extend these ideas to a system with multiple types of customers and servers.

Exploiting Hidden Convexity for Optimal Flow Control in Queueing Networks

C. Bandi, Kellogg School of Management, Northwestern University, c-bandi@kellogg.northwestern.edu
G. G. Loke, NUS Business School, National University of Singapore, gargoeiloke@gmail.com

Optimal flow control in queueing networks is a challenging problem occurring in many contexts, such as data centers, cloud computing, healthcare, revenue management, and distributed networks, etc. The traditional approach has been to adopt heuristic solutions or consider infinite-horizon fluid or diffusion approximations. Motivated by emerging techniques in Robust Optimization, we propose a framework, termed Pipeline Queues, which tracks the dynamics of a queue simultaneously in terms of its queue length and waiting time. We begin by showing that the dynamics of a traditional queueing system can be equivalently modeled using this approach. Our key contribution is the uncovering of the hidden convexity resulting from our modeling approach. This leads us to tractable optimization formulations for generic flow control problems of obtaining performance guarantees on average and quantiles of waiting time, under arbitrary arrival and service distributions with non-zero initial conditions. Our model is flexible enough to capture partial observability and uncertainty of the initial state, as well as various constraints on the control policy. We apply our approach to multiple examples from the literature and numerically illustrate their application. Finally, we implemented our model on a real dataset at a major hospital in India. Our proposed policies are near optimal and perform significantly better than present heuristics.

Robust Queue Inference Engine

C. Bandi, Northwestern University, USA, c-bandi@kellogg.northwestern.edu
A. Proskynitopoulos, Northwestern University, USA, alexej.proskynitopoulos@kellogg.northwestern.edu

We consider the statistical study of partially observed queueing systems arising in application areas such as hospital networks, data centers and cloud computing systems. Since these services operate under strict performance requirements, a statistical understanding of their performance is of great practical interest. A key challenge in these settings is that data is typically incomplete, as recording detailed information on every event in a heavily used system may require substantial overheads and is often not feasible. We propose an analytically tractable framework for studying inference problems in these systems. Our approach is based on the robust queueing framework considered in Bandi et al. (2015), where the queueing primitives are modeled via polyhedral uncertainty sets. Our goal is to infer the uncertainty sets of queueing primitives given partial information. Specifically, we consider the problem of inferring the service distribution given waiting time data and partial information about the arrival process. We construct the uncertainty set for the service process given the information on the waiting times and the uncertainty set for the arrival process. We characterize this set for a variety of queueing systems and demonstrate how uncertainty sets can be leveraged to compute various statistics of the service time. We also report results from an implementation of our framework at one of the biggest hospitals in India, where we estimate the service times at various locations throughout the hospital using partial patient flow data. Our methodology achieves significant computational tractability and provides accurate approximations for primitives of the service process relative to simulated values.

Session 9.2 in Room P3 – Asymptotics and Algorithms on Random Graphs
Chair: Mariana Olvera-Cravioto

Connectivity of a general class of inhomogeneous random digraphs

J. Cao, University of California, Berkeley, jycao@berkeley.edu
M. Olvera-Cravioto, University of North Carolina, Chapel Hill, molvera@email.unc.edu

We study a family of directed random graphs whose arcs are sampled independently of each other, and are present in the graph with a probability that depends on the attributes of the vertices involved. In particular, this family of models includes as special cases the directed versions of the Erdős-Rényi model, graphs with given expected degrees, the generalized random graph, and the Poissonian random graph. We establish the phase transition for the existence of a giant strongly connected component and provide some other basic properties, including the limiting joint distribution of the degrees and the mean number of arcs. In particular, we show that by choosing the joint distribution of the vertex attributes according to a multivariate regularly varying distribution, one can obtain scale-free graphs with arbitrary in-degree/out-degree dependence.

Recursive functions on conditional Galton–Watson trees

N. Broutin, Sorbonne Université, France, nicolas.broutin@upmc.fr
L. Devroye, McGill University, Canada, lucdevroye@gmail.com
N. Fraiman, University of North Carolina at Chapel Hill, USA, fraiman@email.unc.edu

A recursive function on a tree is a function in which each leaf has a given value, and each internal node has a value equal to a function of the number of children, the values of the children, and possibly an explicitly specified random element. The value of the root is the key quantity of interest in general. In this talk, we describe the limit behavior when the leaf values are drawn independently from a fixed distribution, and the tree is a random Galton–Watson tree conditional on its size.

Non-parametric change point detection in growing networks

S. Banerjee, University of North Carolina, Chapel Hill, sayan@email.unc.edu
S. Bhamidi, University of North Carolina, Chapel Hill, bhamidi@email.unc.edu
I. Carmichael, University of North Carolina, Chapel Hill, iain@unc.edu

Motivated by applications of modeling both real world and probabilistic systems such as recursive trees, the last few years have seen an explosion in models for dynamically evolving networks. In this talk, we consider models of growing networks which evolve via new vertices attaching to the pre-existing network according to one attachment function f till the system grows to size $\tau(n) < n$, when new vertices switch their behavior to a different function g till the system reaches size n. We explore the effect of this change point on the evolution and final degree distribution of the network. In particular, we consider two cases, the *standard model* where $\tau(n) = \gamma n$ as well as the *quick big bang model* when $\tau(n) = n^\gamma$ for some $0 < \gamma < 1$. In the former case, we obtain deterministic 'fluid limits' to track the degree evolution in the sup-norm metric and in the latter case, we show that the effect of the pre-change point dynamics 'washes out' when the network reaches size n. We also devise non-parametric, consistent estimators to detect the change point. Our methods exploit and develop new techniques connecting inhomogeneous continuous time branching processes (CTBP) to the evolving networks.

PageRank under degree correlations

M. Olvera-Cravioto, UNC Chapel Hill, United States, molvera@unc.edu

PageRank is one of the most popular algorithms for identifying influential nodes in a directed graph. It has been empirically known for a long time that in scale-free graphs the distribution of the ranks produced by PageRank tends to decay at the same rate as the distribution of the graph in-degree, an observation that has been mathematically proven for a few key random graph models. However, little attention has been given to the effect that significant correlations between the in-degree and out-degree of a vertex can have on the asymptotic behavior of PageRank or other centrality measures. In this talk we present a full characterization of the PageRank distribution for arbitrarily dependent in- and out-degrees in the directed configuration model and the inhomogeneous random digraph. Interestingly, while in the independent case there are two possible ways for a node to be highly ranked (by either having many inbound neighbors or by having a very influential inbound neighbor), in the dependent case the effect of influential neighbors can be completely eliminated, making PageRank a "pure popularity" measure.

Session 9.3 in Room P4 – Queueing and Scheduling
Chair: John Hasenbein

Newsvendor Equations for Production Networks

I. Adan, Eindhoven University of Technology, Netherlands, I.Adan@tue.nl
Z. Atan, Eindhoven University of Technology, Netherlands, Z.Atan@tue.nl
S. W. F. Jansen, Eindhoven University of Technology, Netherlands, S.W.F.Jansen@tue.nl
A. G. de Kok, Eindhoven University of Technology, Netherlands, A.G.d.Kok@tue.nl

We consider production networks with stochastic activity leadtimes. When activities finish early, holding costs are incurred and when end products are delivered late, penalty costs are incurred. The objective is to find the activity start and finish times that minimize the total cost. We introduce the concept of a tardy path and derive the optimality equations for each node in the network. We show that under the optimal solution, for a set of nodes the tardiness probability satisfies the Newsvendor equation.

The Potentially Negative Effects of Cooperation in Service Systems

H. Chung, Korea Advanced Institute of Science and Technology, Korea, hakjin.chung@kaist.ac.kr
H. Ahn, University of Michigan, USA, hsahn@umich.edu
R. Righter, University of California, Berkeley, USA, rrighter@ieor.berkeley.edu

It is well known that the performance of multi-agent service systems degrades with the agents' selfishness (anarchy); this effect is often called the price of anarchy. We investigate a service model in which both customers and the service provider may be strategic. We find that, for a Stackelberg game in which the server invests in capacity before customers decide whether to join or not, there can be a 'Benefit of Anarchy,' i.e., customers acting selfishly can have a greater overall utility than customers who are coordinated to maximize their overall utility. We also show that customer anarchy can be socially beneficial, resulting in a 'Social Benefit of Anarchy.' We show that such phenomena are rather general and can arise in multiple settings (e.g., the firm maximizes its profit vs. the social welfare; the firm chooses capacity vs. price; the service queue is observable vs. unobservable). However, the underlying mechanism leading to the Benefit of Anarchy can significantly differ from one setting to the other.

A New Askew View of c-μ

J. Hasenbein, University of Texas at Austin, USA, jhas@mail.utexas.edu
P. Invernizzi, Politecnico di Milano, Italy, pietrobbe.inve@gmail.com
S. Ripanti, Politecnico di Milano, Italy, ripantisimone@gmail.com

We analyze a multiclass single-server queueing system in which the arrival rates depend on the current job in service. The system is characterized by a matrix of arrival rates. Using fluid models, we obtain the necessary and sufficient conditions for stability. In addition, we are able to characterize the optimal scheduling policy using a Markov decision process representation. The optimal policy is a "twisted" version of the c-μ rule and seems to have some connection with Klimov's rule for standard multiclass, single-server systems.

Fluid models for queueing networks

G. Weiss, The University of Haifa, Israel, gweiss@stat.haifa.ac.il

Fluid models for queueing networks emulate the behavior of the discrete stochastic queueing network by a deterministic continuous fluid network, based on first moment data. They are obtained as limits from uniform scaling of time and space, and can be used to study the transient behavior of the stochastic system. We survey the virtues and limitations of fluid models for verifying stability, and for optimal control of various systems.

Session 9.4 in Room P5 – Jump Markov Processes
Chair: Eugene Feinberg

On Adaptive Control for Continuous-Time Markov Decision Processes

W. Cowan, Rutgers University, USA, c.wes.cowan@gmail.com
D. Ghosh, Rutgers University, USA, dg708@business.rutgers.edu
O. Kanavetas, Rutgers University, USA, okanavetas@business.rutgers.edu
M. Katehakis, Rutgers University, USA, mnk@rutgers.edu
S. Papadimitriou, Rutgers University, USA, spapadim@business.rutgers.edu
D. Pirutinsky, Rutgers University, USA, dp771@scarletmail.rutgers.edu

We consider a Continuous-Time Markov Decision Process (CMDP) with initially unknown transition and/or reward rates. Under a standard irreducibility assumption we develop asymptotically optimal adaptive policies. We discuss applications in healthcare recommendation and scheduling systems.

Non-singular rate matrices in the successive lumping method for Quasi-Skipfree processes

F. M. Spieksma, Universiteit Leiden, The Netherlands, spieksma@math.leidenuniv.nl
D. Ertiningsih, Universitas Gadjah Mada, Yokyakarta, Indonesia, dwi_ertiningsih@ugm.ac.id

A well-known method for the computation of the stationary distribution of a homogeneous Quasi-Birth-Death Process uses the so-called rate matrix, which maps the stationary probabilities of a ‘level’ of the process to the stationary probabilities of the next higher level. We will call this the upward rate matrix. This upward rate matrix is the solution of a quadratic matrix equation. In the case of a Quasi-Skipfree process that is skipfree to the right, such an upward rate matrix exists as well. Unfortunately, it is rather more complicated to compute. However, many applications have additional structure, that can be exploited.
In this talk, we assume that the transition rates to the next higher level having a rank 1 structure (because of the skipfree to the right property, higher levels than the next highest one cannot be reached directly). As is the case for Quasi-Birth-Death processes, the stationary distribution can be shown to have a level-product form. Although in general no easily obtained closed form expression exists for the upward rate matrix, it is simply shown that the factor in the level-product form is the largest eigenvalue of the rate matrix.
We show that a rank 1 structure is equivalent to the existence of a so-called entrance state per level. In the case of the skipfree property to the right, this enables to derive the downward rate matrix (mapping the stationary probabilities of a level to the stationary probabilities of the next *lower* level) explicitly by using ‘successive lumping’. The desired factor in the level product form can be obtained from this explicit rate matrix. Also, under extra invertibility conditions, an alternative form of the more suitable upward rate matrix can be obtained.
We apply our method to some queueing examples.

Kolmogorov’s forward equation for jump Markov processes and its application to control problems

E. A. Feinberg, Stony Brook University, USA, eugene.feinberg@stonybrook.edu
M. Mandava, Indian School of Business, India, Manasa_Mandava@isb.edu
A. N. Shiryaev, Steklov Mathematical Institute and Moscow State University, Russia, albertsh@mi-ras.ru

We describe sufficient conditions for the existence of solutions to Kolmogorov’s forward equation for a given unbounded measurable transition intensity function, sometimes called a Q-function. The related questions are whether a jump Markov process with the given transition intensities exist, is unique, and how its transition probabilities are relevant to the solutions of Kolmogorov’s forward equation. We prove under broad assumptions the existence of the minimal solution, and this solution is the transition probabilities of the corresponding jump Markov process. This problem was studied by William Feller in 1940, and the results there were obtained under narrower conditions than the ones under which the problem was formulated. In fact, Feller’s original formulations hold under the assumptions he considered and even under broader assumptions that take place for controlled jump Markov decision processes (CTMDPs). We use these results to prove that for broad classes of performance criteria, including the expected discounted and undiscounted costs and average costs per unit time, for every policy there exists a Markov policy with the same or better performance vector.

Session 9.5 in Room M1 – Learning in Service Systems
Chair: Stella Kapodistria

Continuous assortment optimization with logit choice probabilities

Y. Peeters, University of Amsterdam, the Netherlands, y.a.peeters@uva.nl
A. den Boer, University of Amsterdam, the Netherlands, a.v.denboer@uva.nl
M. Mandjes, University of Amsterdam, the Netherlands, m.r.h.mandjes@uva.nl

Assortment planning is the field of research dedicated to guiding retailers to make better decisions regarding their assortment. In this field, a widely used model to capture the behavior of customers is the multinomial logit choice model. The utility parameters of this model can be viewed as constants. However, it is realistic to assume that more features play a role in the customers' utility. We regard these features as having a continuous nature. Suppose there are N features that we take into consideration. If these features are controllable, then what is the optimal subset of the N-dimensional feature space that a retailer should offer to its customers? We use a multi-armed bandit approach, that is, we assume a finite selling period in which we have to balance learning the demand and exploiting that knowledge. We present a policy and derive an upper bound on how badly that policy performs. It turns out that any other policy in the worst case performs asymptotically at least as badly as the presented policy.

Managing Ultra-Fast Delivery Service Using Personalized Promotion Strategy for On-line

Z. Saberi, UNSW Canberra, Australia, z.saberi@student.adfa.edu.au
O. Hussain, UNSW Canberra, Australia, o.hussain@adfa.edu.aui
M. Saberi, UNSW Canberra, Australia, m.saberi@unsw.edu.au
E. Chang, UNSW Canberra, Australia, e.chang@adfa.edu.au

E-tailing is one of the main types of retailing which has distinctive feature over its traditional mode, brick, and mortar. One of the features of E-tailing is ultra-fast delivery which is growing fast. However, retailers face a challenge to satisfy this demand in busy days or high traffic zone or time. This paper presents a novel method by shift the demand to more quiet time which allows ultra-fast delivery to happen and keep its customers satisfy. To do this, we leverage the power of personalized promotion which is can be done more effectively now with availability and granularity of large data. Personalized promotion is an efficient tool for individual price discrimination which drives sale up and makes customer relationship stronger. However, we use it to manage our ultra-fast delivery service. To do that, we focus on service promotion rather than price promotion and aim to manage busy delivery day by moving some of the demand to a more quiet time. By accessing to customer purchase history and accordingly offering free delivery service, a customer has more encouragement to use this opportunity to move their shopping events. Not only customer enjoy free delivery, but the e- tailer can also manage the smoothness routine of delivery. The proposed model comprises three tasks. In Task 1, the delivery load of each week is predicted using both historical and external data (traffic rate) is predicted. In Task 2, we measure the utility of customer in buying a given item one or two days before the forecasted heavy delivery day using multiple discrete choice model. Task 3 is devoted to identifying customers who may accept our service promotion offer by using heuristic optimization model.

Decision making under uncertainty

S. Kapodistria, Eindhoven University of Technology, The Netherlands, s.kapodistria@tue.nl

We consider a class of stochastic dynamic programming (SDP) problems arising in the context of maintenance. Such SDP problems aim at deriving optimal policies (when should maintenance be performed and by how much should the system be maintained). From a fundamental perspective, the difficulty of such SDP problems lies in deriving the optimal policy. From a numerical perspective, the difficulty lies in computing the policy (the curse of dimensionality). While, from a practical perspective, the difficulty lies in the stochastic modelling of all the relevant information (big data). In this talk, we present an overarching framework that aims at addressing all aforementioned challenges.

Session 9.6 in Room M2 – Random Walks/Dynamic Markets
Chair: Ruiting Zuo

Conditions for recurrence and transience for one family of random walks

V. M. Abramov, Monash University, Australia, vabramov126@gmail.com

In the present talk, a family of two-dimensional random walks $\{\mathbf{S}_t, \mathcal{A}\}$ in the main quarter plane, where $\mathcal{A}$ is a set of infinite sequences of real values, is discussed. For $a \in \mathcal{A}$, a random walk is denoted $\mathbf{S}_t(a) = (S_t^{(1)}(a), S_t^{(2)}(a))$. Let θ denote the infinite sequence of zeros. For $a \neq \theta$ the components $S_t^{(1)}(a)$ and $S_t^{(2)}(a)$ are assumed to be correlated in the specified way that will be defined exactly in the talk, while for $a = \theta$, the random walk $\mathbf{S}_t(\theta)$ is the simple two-dimensional random walk in the main quarter plane. We find the conditions on a under which a random walk $\mathbf{S}_t(a)$ is recurrent or transient. In addition, we introduce new classes of random walks, ψ-random walks, and find conditions under which a subfamily of random walks $\{\mathbf{S}_t, \mathcal{A}_\psi\}$, $\mathcal{A}_\psi \subset \mathcal{A}$ belongs to the class of ψ-random walks.

Geometric properties of moment generating functions of increments for a three dimensional reflecting random walk

M. Kobayashi, Tokai University, Japan, m_kobayashi@tsc.u-tokai.ac.jp

We are concerned with a stationary distribution of a k-dimensional reflecting random walk, where k is a natural number which represents a dimension. For $k = 1$, it is well known that the stationary distribution is given by the geometric distribution. However, for $k \geq 2$, it is not expected that the stationary distribution has a closed form expression. Thus, we will obtain results in another directions. In particular, for $k = 2$, the stability condition, tail asymptotics, product form condition, etc. are verified by many papers (see, e.g., Fayolle, Malyshev, and Menshikov (1995), Kobayashi and Miyazawa (2013), Ozawa (2013), Latouche and Miyazawa (2014), Ozawa and Kobayashi (2018), and references in those papers), but for $k \geq 3$, there are few theoretical results. Hence, we are interested in deriving them under the case $k \geq 3$. For $k = 2$, the convergence domain of moment generating function for the stationary distribution has a key roll to obtain the stability condition, tail asymptotics, product form condition. In addition, its convergence domain is determined by geometric properties of moment generating functions of increments. That is, their geometric properties are closely related with the stationary distribution. In this talk, for $k = 3$, we introduce their relationships and give the light-tailed and product form conditions by the geometric properties.

Dynamic Market-Making: A Myersonian Approach

M. Akbarpour, Stanford University, United States, mohamwad@stanford.edu
E. V. Muir, Stanford University, United States, evmuir@stanford.edu
P. Strack, UC Berkeley, United States, philipp.strack@gmail.com

We study a dynamic mechanism design environment in which units of a homogeneous good are traded among impatient buyers and sellers with private types that arrive according to independent Poisson processes. The mechanism designer acts as a broker that accumulates inventory by purchasing the good from arriving sellers in order to meet demand from future buyers. Under the optimal allocation, a unit of the good is sold to an arriving buyer if and only if the buyer's virtual valuation exceeds the cost of losing a unit of inventory (which decreases the continuation value of the designer). Similarly, a unit of the good is purchased from an arriving seller if and only if the seller's virtual cost is less than the benefit associated with a one unit increase in inventory (which increases the continuation value of the designer). Thus, the profit-maximizing mechanism can be implemented via a simple posted-price mechanism with bid and ask prices that adjust purely as a function of the inventory of the designer. We also consider the implications of this mechanism for the inventory accumulation process (which is a simple birth-and-death process), the bid-ask spread and markup charged by the designer, price volatility, market thickness and competition among market-markers.

Dynamic Contracting in Worker-Manager-Owner Relationship

J. Keppo, National University of Singapore, Singapore, keppo@nus.edu.sg
N. Touzi, Ecole Polytechnique, France, nizar.touzi@polytechnique.edu
R. Zuo, National University of Singapore, Singapore, ruitingz@u.nus.edu

In typical contracting problems, a principal designs a contract and then an agent decides whether to accept it and take corresponding actions or to reject it. In this paper, we investigate incentive contracts under dynamic actions and a simple organizational structure with an owner of an organization, a manager, and a worker. The owner cannot perfectly observe the manager and worker's actions, and similarly, the manager cannot perfectly observe the worker's actions. We show how the actions of the participants and the costs of their actions interact.

Session 9.7 in Rooms M5 and M6 – Analysis of Computing and Service Systems
Chair: Sherwin Doroudi

Scalable Load Balancing in the Presence of Servers Under Interference

J. Abdul Jaleel, University of Minnesota, USA, abdul314@umn.edu
S. Doroudi, University of Minnesota, USA, sdoroudi@umn.edu
M. Delasay, Stony Brook University, USA, mohammad.delasay@stonybrook.edu

Virtualization is a common feature in modern cloud data centers that allows each physical server to host multiple virtual machines (VMs), not all of which are necessarily employed by the same application. Despite the many benefits offered by virtualization, not all physical resources are easy to partition across virtual machines located on the same physical server. This can lead to unpredictable instances of temporary performance degradation, known as interference. We consider the load balancing problem encountered by a cloud application hosted across several VMs (located on different physical machines). The application faces a stream of requests that must immediately be dispatched to one of the VMs with the goal of minimizing average response times (latency), a problem that is further complicated by the fact that each VM may be experiencing unpredictable (but detectable) interference, which we model as a Markov-modulated process. We propose a number of distributed dispatching policies that dispatch incoming requests based on the interference and/or busy status of a randomly sampled subset of the servers. Using a combination of mean field analysis and recursive renewal reward methods, we evaluate the performance of these heuristics in a variety of settings, while deducing a number of surprising results.

An Exact Analysis of a Class of Markovian Bitcoin Models

K. Javier, Clemson University, USA, kjavier@clemson.edu
B. Fralix, Clemson University, USA, bfralix@clemson.edu

We consider two different continuous-time Markov chain models recently studied in Göbel et al. (2016), which were created to model the interactions between a small pool of miners, and a larger collection of miners, within the Bitcoin blockchain. The first model we discuss, which we refer to as Model 1, represents the case where all miners behave honestly and follow the Bitcoin protocol, while the second model, Model 2, represents the case where a pool of miners use the selfish-mining strategy. We give a new derivation of the stationary distribution of Model 1 and further build on the results of Göbel et al. by showing that the normalizing constant can be expressed in closed-form. We also illustrate how these techniques yield similar expressions for the stationary distribution of Model 2 as well, which can be calculated exactly using a finite number of operations. Results that address the time-dependent (transient) behavior of these models, as well as possible generalizations, will be discussed as time permits. Our methods make use of the recently-developed random-product technique introduced in Buckingham et al. (2015).

Dispatching in Heterogeneous Systems

K. Gardner, Amherst College, kgardner@amherst.edu
C. Stephens, Amherst College, cstephens19@amherst.edu

In multiserver distributed systems, the question of how to dispatch an arriving job to a server is key in reducing response time. When service times are exponential and all servers work at the same rate, the classical Join-the-Shortest-Queue (JSQ) policy is known to be optimal. However, large scale systems typically consist of heterogeneous servers; when some servers are faster than others, policies like JSQ are no longer optimal. Instead, we may prefer a policy like Shortest Expected Delay (SED), which dispatches an arriving job to the server at which its expected response time (based on the number of jobs in the queue and the server's speed) is the shortest. The problem is even further complicated if we do not know the relative speeds of different servers, or even which servers are fast and which are slow. In this work, we propose and analyze new dispatching policies for heterogeneous systems with limited information about server speeds. We find that in many cases, simple heuristics yield performance as good as or even better than the performance under classical dispatching policies.

Dynamic Market Prices for Call Centers with Co-Sourcing Subject to Non-stationary Demand

S. Doroudi, University of Minnesota, U.S., sdoroudi@umn.edu
M. Mousavi, Target Corp., U.S., ehsan.mousavi@gmail.com

We consider a call center that can offload (or co-source) some or all of its arriving customers to an external provider. These customers are impatient—in the sense that they will eventually abandon the queue if left unserved—and arrive according to a non-stationary but known arrival pattern during a finite time horizon. By co-sourcing customers, the call center can reduce delays and curtail customer abandonment, however for each customer the call center co-sources, the call center must pay the external provider a price based on the time of transfer. We view the resulting queueing system as a fluid model where the time horizon naturally separates into alternating overloaded, critically loaded, and underloaded time periods. We characterize the optimal proportion of co-sourced demand and the equilibrium co-sourcing prices, both as a functions of time. We find the existence of "all or nothing" optimal co-sourcing proportions: at any given point in time—unless the sytem is under critical load—either all arriving customers should be co-sourced, or none of them should be co-sourced (i.e., they should all be served in-house). Moreover, we find that both extremes can exist in both overloaded and underloaded time intervals.

Session 9.8 in Rooms M7 and M8 – Stochastic Modeling for Service Systems
Chair: Nilay Tanik Argon & Lerzan Örmeci

Halfin-Whitt scaling of the Fixed-Cycle Traffic-Light queue

M. A. A. Boon, Eindhoven University of Technology, The Netherlands, m.a.a.boon@tue.nl
A. J. E. M. Janssen, Eindhoven University of Technology, The Netherlands, a.j.e.m.janssen@tue.nl
J. S. H. van Leeuwaarden, Eindhoven University of Technology, The Netherlands, j.s.h.v.leeuwaarden@tue.nl
R. W. Timmerman, Eindhoven University of Technology, The Netherlands, r.w.timmerman@tue.nl

The Fixed-Cycle Traffic-Light queue is the standard model for intersections with static signaling. Vehicles arrive at a lane at the intersection, controlled by a traffic light, form a queue and depart during fixed cycles. Recently, a Pollaczek contour-integral expression for the generating function of the queue length in a single lane has been derived. This result can be used to obtain results on the queue-length distribution in e.g. the Quality-and-Efficiency Driven (QED) regime, when choosing a Halfin-Whitt scaling. The QED regime ensures (a.o.) that, even when the load on the system tends to 1, there is a strictly positive probability that arriving vehicles experience no delay.
Moreover, in the QED regime the performance measures, like the mean queue length, have nice forms. In a setting where the green and red times have to be chosen for a whole intersection, these nice forms can be used to find optimal settings for the traffic light. The used scaling, and therefore part of the constraints of the optimization problem, remind of classical staffing rules in e.g. call centers. Our results are more generally applicable, for example also to dimensioning problems in e.g. the bulk-service queue and relaxations of the Fixed-Cycle Traffic-Light queues.

Assigning Priorities (or not) in Service Systems with Nonlinear Waiting Costs

H. Ouyang, The University of Hong Kong, Hong Kong, oyhy@hku.hk
N. Argon, University of North Carolina at Chapel Hill, United States, nilay@email.unc.edu
S. Ziya, University of North Carolina at Chapel Hill, United States, ziya@email.unc.edu

For a queueing system with multiple types of customers differing in service time distributions and costs for waiting, it is well known that the $c\mu$ rule is optimal if costs for waiting are incurred linearly with time. It is also known that when costs are convex increasing functions, the generalized $c\mu$ rule, which is dependent on queued customers' already experienced waiting times, is asymptotically optimal under heavy traffic. In this paper, we seek to identify policies that minimize the long-run average cost under possibly nonlinear waiting cost functions and arbitrary traffic conditions but within the set of static policies that require information only on the type identities of jobs and their order of arrival. For a single-server queueing system with two types of customers, our main result gives conditions under which type-based priority policies and the first-come-first-serve (FCFS) policy can be ordered for non-decreasing cost functions that are first-order differentiable. We then apply this result to polynomial cost functions and obtain useful insights into when prioritization should be preferred over FCFS and when it should be avoided. For example, unlike in the linear-cost case, it could be better not to give priority to a certain type at all and employ FCFS under quadratic cost functions, especially when the traffic intensity is high. Finally, by means of a numerical study, we test how the best static policy compares with the generalized $c\mu$ rule, which requires information on the current waiting times of customers and precise structure for the cost functions. In these experiments with quadratic waiting costs, we find that using the best static policy (type-based priority or FCFS) performs comparably with (sometimes even better than) the generalized $c\mu$ rule except when the traffic intensity is high and there is not a clearly more “important” type in regards to dominance with respect to rates of cost and service.

When to Triage in Service Systems with Hidden Customer Class Identities?

Z. Sun, City University of Hong Kong, Hong Kong, zhankun.sun@cityu.edu.hk
N. T. Argon, University of North Carolina at Chapel Hill, USA, nilay@unc.edu
S. Ziya, University of North Carolina at Chapel Hill, USA, ziya@unc.edu

In service systems with heterogeneous customers, prioritization with respect to the relative importance of customers is known to improve certain performance measures. However, in many applications, information necessary to determine the importance level of a customer may not be available immediately but can be revealed only through some preliminary investigation, referred to as *triage*. This triage process is typically error-prone and may take substantial amount of time, and hence, it is not always clear if and when it should be implemented for purposes of priority assignment. To provide insights into this question, we study a stylized queueing model with a single server and two types of customers with hidden type identities and different rates of waiting costs. We first show that the optimal dynamic policy on triage is characterized by a switching curve: if the number of unclassified customers is sufficiently large, then the server triages an unclassified customer; otherwise, the server skips triage and directly serves an unclassified customer. We also identify a mathematical condition under which triage should never be considered. Based on all our numerical and analytical results, we conclude that triage for purposes of prioritization should not be considered if it takes a significantly large amount of time in comparison to the actual service but if it is sufficiently fast, then implementing a policy that triages some or all customers can significantly reduce long-run average costs of waiting. In these cases where there is benefit to using triage, a policy that triages all arrivals can be used if the triage time is substantially short and the probability that a customer is classified into a higher priority is large; otherwise, a policy that triages customers dynamically depending on the system state would perform significantly better than the triage-all policy, especially when the traffic is heavy.

Strategic Customers in Systems with Batch Arrivals

L. Örmeci, Koç University, Turkey, lormeci@ku.edu.tr
O. Bountali, University of Toronto, Canada, Olga.Bountali@Rotman.Utoronto.Ca
A. Burnetas, University of Athens, Greece, aburnetas@math.uoa.gr

Customers who arrive in groups at a service or production system but are served individually are faced with a convoluted manifestation of the “join-balk” dilemma when they try to balance the payoff obtained from service with the costs caused by waiting in the queue: if they are bound to join or balk as a group, then some of them may be forced to join so that the net benefit of the entire group is maximized. On the other hand, if separate customer entrances are allowed, then their individual interest plays the major role in their decision. In this sense, when all arriving groups act strategically, it is of interest to predict the resulting join- balk strategies under equilibrium in both cases, as well as their differences in system performance, customer throughput, total customer welfare, etc. We consider these questions in the framework of a single server Markovian queue with batch arrivals and random batch sizes. We explicitly consider two cases with respect to the entrance decision rules: the 0-1 case, in which the entire batch (i.e., all its customers) decides to either join or balk, and the partial case, where the batch is allowed to join partially (i.e., only a fraction of the customers decides to join). We analyze and compare how the customers behave in equilibrium under both rules and the corresponding implications on the social welfare.

9

Author Index

X

Y

Z

Booklet compiled by:

Thomas Taimre The University of Queensland, Australia

Marijn Jansen The University of Queensland, Australia

Thanks to Marko Boon (Eindhoven University of Technology, the Netherlands) for developing the software to create this booklet and to the UQ Images team for permission to use the photograph on the cover.

This version: July 1, 2019