2016 Stochastic Networks Conference
Poster Session Abstracts

June 20, 2016

1. Agarwal, Pooja
   Division of Applied Mathematics, Brown University
   Equilibria of randomized load balancing algorithms with general service distributions
   pooja_agarwal@brown.edu

2. Aghajani, Reza
   Division of Applied Mathematics, Brown University
   Mean-Field Dynamics of Load-Balancing Networks with General Service Distributions
   mohammadreza_aghajani@brown.edu

3. Aksoy, Sinan
   Mathematics Department, University of California San Diego
   Extreme values of the stationary distribution of random walks on directed graphs
   saksoy@ucsd.edu

4. Barrera, Javiera
   School of Engineering and Sciences, Universidad Adolfo Ibáñez
   Calibration of a dependent failure model and the topological optimization of reliable network
   javiera.barrera@uai.cl

5. Bayati, Mohsen
   Stanford University
   Online Decision-Making with High-Dimensional Covariates
   bayati@stanford.edu

6. Carmen, Raisa
   Department of Decision Sciences and Information Management, KU Leuven
   A queueing model to analyse the impact of boarding in the emergency department
   raisa.carmen@kuleuven.be

7. Cecchi, Fabio
   Eindhoven University of Technology
   Mean-Field Analysis of large-scale random medium access algorithms
   F.Cecchi@tue.nl

8. Costantini, Cristina
   Dipartimento di Economia, Universita' di Chieti-Pescara
   Well posedness of constrained martingale problems for reflecting diffusions in piecewise smooth domains
   c.costantini@unich.it

9. Feng, Aurora Jiekun
   Department of Statistical Science, Cornell University
   Steady-state Diffusion Approximations for Discrete-time Queue in Hospital Inpatient Flow Management
   jf646@cornell.edu

10. Ferragut, Andrés
    MATE Research Group; Universidad ORT Uruguay
    Optimal timer-based caching policies under general heavy-tailed request processes
    ferragut@ort.edu.uy

11. Friedlander, Eric
    Department of Statistics and Operations Research, University of North Carolina at Chapel Hill
    Diffusion Approximations for Controlled Weakly Interacting Systems
    ericf2218@gmail.com

12. Gerencsér Balázs
    Probability & Statistics research division, Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences
    Robust averaging - performance of the push-sum algorithm in the presence of transmission failures
    gerencser.balazs@renyi.mta.hu

13. Hermansson, Niffe
    Department of Statistics, University of Auckland
    User equilibria in parallel Processor Sharing queues
    nher257@aucklanduni.ac.nz

14. Kamphorst, Bart
    Stochastic Department, Centrum Wiskunde & Informatica
    Achievable Performance of Blind Policies in Heavy Traffic
    b.kamphorst@cwi.nl

15. Khezeli, Ali
    Department of Mathematics, Sharif University of Technology
    Stable Transport Between Stationary Random Measures
    alikhezeli@gmail.com

16. Lipshutz, David
    Division of Applied Mathematics, Brown University
    Pathwise differentiability of reflected diffusions
    david_lipshutz@brown.edu

17. Lyu, Hanbaek
    Department of Mathematics, The Ohio State University
    Synchronization of finite-state pulse-coupled oscillators
    colourgraph@gmail.com

18. Manjrekar Mayank
    Department of Mathematics, University of Texas at Austin
    Spatial processes with births and deaths - Hard-core regime
    mayankm@utexas.edu

19. Mukherjee, Debankur
Department of Mathematics and Computer Science, Eindhoven University of Technology
Universality of Power-of-d Load Balancing in Many-Server Systems
d.mukherjee@tue.nl

20. Nesti, Tommaso
Stochastics Group, Centrum Wiskunde & Informatica
Reliability of energy networks under uncertainty: a large deviations approach
T.Nesti@cwi.nl

21. O'Reilly, Elizabeth
Department of Mathematics, University of Texas at Austin
Optimization of DNA sequencing using Stochastic Geometry
eoreilly@math.utexas.edu

22. Palowitch, John
Statistics and Operations Research, University of North Carolina at Chapel Hill
The Continuous Configuration Model: A Null for Community Detection on Weighted Networks
palojj@email.unc.edu

23. Patch, Brendan
The University of Queensland/University of Amsterdam
Detecting Markov Chain Instability: A Monte Carlo Approach
b.patch@uq.edu.au

24. Pender, Jamol
School of Operations Research and Information Engineering, Cornell University
Strong Approximations for Time Varying Queues with Non-Renewal Arrival and Service Processes
jamol.pender@gmail.com

25. Rahimian, Mohammad
Electrical and Systems Engineering, University of Pennsylvania
Moment-Based Spectral Analysis of Random Graphs with Given Expected Degrees
rahimian.amin@gmail.com

26. Reiman, Martin
Industrial Engineering and Operations Research, Columbia University
A Stochastic Programming Based Approach to Control of Assemble-to-Order Inventory Systems
martyreiman@gmail.com

27. Rhee, Chang-Han
Centrum Wiskunde & Informatica
Sensitivity analysis for Markov chains
C.Rhee@cwi.nl

28. Rowat, Peter
Institute for Neural Computation, University of California San Diego
Stochastic network thinking applied to firing patterns of stellate neurons
peter@pelican.ucsd.edu

29. Saha, Subhamay
Department of Electrical Engineering, Technion - Israel Institute of Technology
Optimality of the Generalized cRule in the Moderate Deviation Regime
subhamay585@gmail.com

30. Sloothaak, Fiona
Mathematics and Computer Science, Eindhoven University of Technology
Asymptotic analysis of a cascading failure model
f.sloothaak@tue.nl

31. Uribe, Cesar
Coordinated Science Laboratory, University of Illinois at Urbana-Champaign
Convergence Rates in Distributed Learning: Acceleration, Network Independence and Uniform Social Sampling
cauribe2@illinois.edu

32. Wang, Cheng-Heng
Department of Electrical and Computer Engineering, University of California, San Diego
Adaptive Policies for Scheduling with Reconfiguration Delay: An End-to-End Solution for All-Optical Data Centers *hw009@eng.ucsd.edu*

33. Wang, Alex
Department of Statistics, University of Auckland
Selfish routing in a network of parallel queues
ywan925@aucklanduni.ac.nz

34. Yu, Yao
Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University
Optimal Routing to Remote Queues *yyu15@ncsu.edu*

35. Zocca, Alessandro
Centrum Wiskunde & Informatica
Minimizing heat dissipation in DC networks using batteries
A.Zocca@cwi.nl

36. Zubeldi, Martin
Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology
Delay, memory and messaging tradeoffs in distributed service systems
zubeldia@mit.edu

Stochastic Networks 2016 Poster Application

Title of poster
Equilibria of randomized load balancing algorithms with general service distributions

Name of presenter
Pooja Agarwal

Department and Institution of presenter
Division of Applied Mathematics, Brown University

Abstract
Randomized load balancing algorithms play an important role in large-scale networks. The hydrodynamic or fluid limit of the state of these networks, as the number of nodes goes to infinity, can often be described by systems of (deterministic) coupled measure-valued equations. Under fairly general conditions, we prove the existence of a unique equilibrium point for the hydrodynamic equations associated with several routing algorithms, including the join-the-shortest-of-d-queues and the join-the-idle queue policies. We also discuss some properties of these equilibrium points, and their implications for the performance of these policies.

Mean-Field Dynamics of Load-Balancing Networks with General Service Distributions

Reza Aghajani
Joint work with Kavita Ramanan

Division of Applied Mathematics, Brown University

Abstract. We introduce a general framework for studying a class of randomized load balancing models in a system with a large number of servers that have generally distributed service times and use a first-come-first serve policy within each queue. Under fairly general conditions, we use an interacting measure-valued process representation to obtain hydrodynamics limits for these models, and establish a propagation of chaos result. Furthermore, we present a set of partial integro-differential equations (PDEs) whose solution can be used to approximate the transient behavior of such systems. We prove that these PDEs have a unique solution, use a numerical scheme to solve them, and demonstrate the efficacy of these approximations using Monte Carlo simulations. We also illustrate how the PDE can be used to gain insight into network performance.

Stochastic Networks 2016 Poster: Sinan Aksoy

February 5, 2016

Title of poster: Extreme values of the stationary distribution of random walks on directed graphs

Name of presenter: Sinan Aksoy

Department and Institution of presenter: Mathematics Department, UC San Diego

Short abstract for the poster (300 words or less): We examine the stationary distribution of random walks on directed graphs. In particular, we focus on the *principal ratio*, which is the ratio of maximum to minimum values of vertices in the stationary distribution. We give an upper bound for this ratio over all strongly connected graphs on n vertices. We characterize all graphs achieving the upper bound and we give explicit constructions for these extremal graphs. This is joint work with Fan Chung and Xing Peng.

Calibration of a dependent failure model and topological optimization of reliable networks

Javiera Barrera
School of Engineering and Sciences, Universidad Adolfo Ibáñez,Chile

10 de junio de 2016

Abstract: We address the design problem of a reliable network. Previous works assume that link failures are independent. We discuss the impact of dropping this assumption. We show that under a common-cause failure model, dependencies between failures can affect the optimal design. We consider a graph with a set of vertex and a set of potential links. To do the network design two problems arise, choose and calibrate a dependent failure model and, once we have it, then to choose the optimal set of links to attain the most reliable network. In [4] we propose a calibration for the classical Marshall-Olkin model, we study the empirical evidence that shows that dependence arises in different kind of network but with similar behaviours see [2], [3] and [6]. Based on this evidence we propose a sequential linear programming formulations that allow us to calibrate the parameters of the Marshall-Olkin given the failure probabilities and the failure-correlation between links. We use simulation to test our methodology with a geographical failure model, where failures in links are naturally correlated. For the topological design part, given a calibrated Marshall-Olkin model, in [1] we provide an integer-programming formulation to solve to maximize the reliability. We used Sample Average technique obtaining a similar formulation as the one in [5].

Reference:

1. J.Barrera, H. Cancela and E. Moreno, *Topological optimization of reliable networks under dependent failures*. Operation Research Letters 43 (2), 2015, 132–136.
2. Gonzalez, A. J., B. E. Helvik, J. K. Hellan, and P. Kuusela. 2010. *Analysis of Dependencies between Failures in the UNINETT IP Backbone Network*. In 2010 IEEE 16th Pacific Rim International Symposium on Dependable Computing, 149–156. IEEE.
3. Gill, P., N. Jain, and N. Nagappan. 2011. *Understanding network failures in data centers*. In ACM SIGCOMM Computer Communication Review, Volume 41, 350. ACM.
4. O. Matus, J. Barrera, E. Moreno and G. Rubino *Calibrating A Dependent Failure Model for Computing Reliabilities on Telecommunication Networks* Accepted for WSC 2016.
5. Y. Song, J. Luedtke, Branch-and-cut approaches for chance-constrained formulations of reliable network design problems, Mathematical Programming Computation 5 (4), 2013, 397?432.
6. Turner, D., K. Levchenko, A. C. Snoeren, and S. Savage. 2010. *California fault lines: understanding the causes and impact of network failures*. In ACM SIGCOMM Computer Communication Review, Volume 40, 315–326. ACM.

In colaboration with Héctor Cancela, Universidad de la República, Uruguay. Omar Matus, Universidad Adolfo Ibáñez, Chile. Eduardo Moreno, Universidad Adolfo Ibáñez, Chile. Gerardo Rubino, INRIA Rennes, France.

Online Decision-Making with High-Dimensional Covariates

Mohsen Bayati

Stanford University

Abstract. Large data sets have enabled decision-makers to tailor choices at the individual-level. This involves learning a model of decision rewards conditional on individual-specific covariates. In domains such as medical decision-making and personalized advertising, these covariates are often high-dimensional; however, typically only a small subset of these observed features are predictive of each decisions success. We formulate this problem as a multi-armed bandit with high-dimensional covariates, and present a new efficient bandit algorithm based on the LASSO estimator. Our regret analysis establishes that our algorithm achieves near-optimal performance in comparison to an oracle that knows all the problem parameters. The key step in our analysis is proving a new oracle inequality that guarantees the convergence of the LASSO estimator despite the non-i.i.d. data induced by the bandit policy. Furthermore, we illustrate the practical relevance of our algorithm by evaluating it on a real-world clinical problem of warfarin dosing. A patients optimal warfarin dosage depends on the patients genetic profile and medical records; incorrect initial dosage may result in adverse consequences such as stroke or bleeding. We show that our algorithm outperforms existing bandit methods as well as physicians to correctly dose a majority of patients.

The work is based on [1].

References

[1] H. Bastani and M. Bayati, *Online Decision Making with High Dimensional Covariates*, Preprint available online.

A queueing model to analyse the impact of boarding in the emergency department

R. Carmen

KU Leuven, Department of Decision Sciences and Information Management, {Raisa.Carmen}@kuleuven.be

Abstract

Boarding patients are patients that require admission to the hospital after being treated in the emergency department (ED) but are stranded in the ED because of a lack of beds in the hospital wards. This *'inpatient boarding'* phenomenon is considered to be a big problem in many EDs all over the world and has been associated with increased ambulance diversions, worse patient outcomes, frustration among medical staff, higher patient length of stay, loss of revenue, and higher mortality rates.

The queueing network we are considering, models the ED as a semi-open queueing network with a limited number of beds and physicians. Patients may have to visit the physician more than once and boarding patients impact the treatment process by occupying beds while they wait for admission, preventing newly arriving patients from entering the ED. We analyse and solve our queueing network in an exact numerical way using a Markov-Modulated Fluid Queue (MMFQ). The advantage of the MMFQ over the standard QBD approach for exact analysis, is that service levels (the probability that the waiting time to obtain a bed is smaller than a certain threshold) are obtained more efficiently. We observe that boarding patients can put a lot of pressure on the ED when the number of beds is limited and investigate policies aimed at reducing the number of boarding patients.

Keywords: markov-modulated fluid queue, emergency department, capacity planning

Fabio Cecchi – Eindhoven University of Technology
Address: Eindhoven University of Technology,
P.O. Box 513, 5600 MB Eindhoven, NL
E-mail: f.cecchi@tue.nl
Tel: +31 6 48482723

5th February, 2016

TITLE: *Mean-Field Analysis of large-scale random medium access algorithms*

ABSTRACT:

With the rapid increase in wireless applications in recent years, there has been increasing demand for wireless spectrum and with it the need to develop efficient wireless protocols. Even today wireless networks are already large and complex, and they are to get denser and larger in the future. The development of such networks is at the heart of the so-called Internet of Things' network (IoT). Dedicated medium cannot be assigned uniquely to each node in the network, so that the medium has to be shared by various nodes. Multiple simultaneous node transmissions on the same channel will inevitably give rise to interference and loss of throughput, hereby the necessity to investigate the performance of MAC protocol.
Distributed random MAC algorithms such as CSMA provide a popular mechanism for sharing the transmission medium among competing users in large-scale wireless networks. Conventional models for CSMA that are amenable for analysis assume that users have always packets to transmit. In contrast, when users do not compete for medium access when their buffers are empty, a complex interaction arises between the activity states and the buffer contents. We develop a mean-field approach to investigate this dynamic interaction for networks with many users. We identify a time-scale separation between the evolution of the activity states and the buffer contents, and obtain a deterministic dynamical system describing the network dynamics on a macroscopic scale. The fixed point of the dynamical system yields highly accurate approximations for the stationary distribution of the buffer contents and packet delay, even when the number of users is relatively moderate.
The developed mean-field approach proved to be generalizable to more involved CSMA-like networks, and such extensions are outlined.
(Joint work with S.C. Borst, J.S.H. van Leeuwaarden, P.A. Whiting)

Title of poster: Well posedness of constrained martingale problems for reflecting diffusions in piecewise smooth domains.

Presenter: Cristina Costantini

Affiliation of presenter: Dipartimento di Economia, Universita' di Chieti-Pescara

Abstract:

A recent work by Kang and Ramanan (arXiv:1412.0729, 2014) establishes the equivalence, under the assumption of strict uniform ellipticity, of uniqueness of the solution for stochastic differential equations with reflection and submartingale problems for reflecting diffusions. Combined with the results of Dupuis and Ishii (Ann Probab, 1993) this ensures uniqueness of the solution for a large class of submartingale problems for reflecting diffusions in piecewise smooth domains, a key tool in diffusion approximation of queueing networks.

Considering constrained martingale problems, uniqueness of the solution can be obtained, for the equivalent class of problems, by Costantini and Kurtz (EJP, 2015) and the comparison principle proved in Dupuis and Ishii (Hokkaido Math J, 1991), without the assumption of strict uniform ellipticity.

This allows to deal with reflecting diffusion processes with some components that follow deterministic dynamics (given the other components), as might arise, for instance, in a multiscale limit of a network.

The result carries over to sticky reflection diffusion processes as well, at least under some assumptions.

Our approach can be employed also for infinite-dimensional processes, for instance measure valued processes.

Joint work with T.G. Kurtz

Steady-state Diffusion Approximations for Discrete-time Queue in Hospital Inpatient Flow Management

Aurora Jiekun Feng
Department of Statistical Science, Cornell University
Joint work with Pengyi Shi
Krannert School of Management, Purdue University

Motivated by the recent development of steady-state approximation via Stein's method for Erlang-A and Erlang-C models, we apply the framework to a discrete-time Markov chain (DTMC) setting. This DTMC is motivated from studying hospital inpatient flow. It captures the patient midnight census in the inpatient department, which is a key performance metric monitored by many hospitals and is essential in predicting the patient census during different times of a day. Specifically, the DTMC $\{X_k : k = 0, 1, \dots\}$ is characterized by $X_{k+1} = X_k + A_k - D_k$, where A_k and D_k denote the total number of patient arrivals and discharges within day k, respectively. A_k follows an i.i.d Poisson distribution with a predetermined, fixed rate, and D_k follows a binomial distribution with parameters depending on X_k and service rate μ. Using Stein's method, we derive a constant-variance diffusion process to approximate the stationary distribution of a scaled DTMC $\{\tilde{X}_k : k = 0, 1, \dots\}$ defined as $\tilde{X}_k = 1/\sqrt{R}(X_k - R)$, with R denoting the offered load. Different from the continuous-time counterpart, i.e., the Erlang-C model, we find that the service rate μ plays a critical role in the convergence between the stationary distribution of the diffusion process and that of the scaled DTMC $\{\tilde{X}_k\}$. We prove that, when μ approaches 0 with a rate $1/R$, the Wasserstein distances between the stationary distribution of the diffusion process and that of $\tilde{X}_k$ converges at a rate of $1/\sqrt{R}$. This conclusion is "universal", in the sense that it is valid in any load condition from lightly loaded (quality-driven) to heavily loaded (non-degenerate slowdown). Through numerical experiments, we demonstrate the impact of μ on the quality of the steady-state diffusion approximation. To further improve the approximation quality, we devise another diffusion process whose variance depends on its particular state. We show numerically that this new diffusion approximation can outperform the constant-variance one.

Optimal timer-based caching policies under general heavy-tailed request processes.

Presenter: Prof. Andrés Ferragut
MATE Research Group
Universidad ORT Uruguay

Joint work with I. Rodríguez and F. Paganini

Abstract

In content distribution networks, content caching is the usual way to provide service to end users, helping to reduce service latency and bandwidth usage across the network. Traditionally, the caching algorithms (i.e. how to decide which content must be stored at a given time due to memory constraints) can be split between *replacement policies* and *timer based policies*. The latter have received considerable attention recently because its mathematical analysis is well suited to the theory of point processes.

In this work, we analyze the hit performance of cache systems that receive file requests with general interarrival times and different popularities. We consider timer-based (TTL) policies, with differentiated timers over which we optimize. We formulate an optimization problem to characterize the optimal policy. Interestingly, the convexity structure of this problem depends heavily on the monotonicity of the *hazard rate function* of the interarrival distribution. In particular for decreasing hazard rates (associated with heavy tails), the problem is convex. By analyzing the resulting optimal policy, we show that it outperforms the static policy of caching only the most popular contents.

We provide explicit solutions for the optimal policy in the case of Pareto-distributed inter-request times and Zipf distribution of file popularities, including a compact fluid characterization in the scaling limit of a large number of files. Through simulations, we compare the performance against classical replacement policies, such as least-recently-used. Finally, we analyze extensions of the optimization framework to line networks of timer based caches.

Diffusion Approximations for Controlled Weakly Interacting Systems

Eric Friedlander
Department of Statistics and Operations Research
University of North Carolina at Chapel Hill

Abstract

We consider a rate control problem with a discounted cost criterion for an N-particle, weakly interacting, pure jump, finite state Markov process. Such models arise from problems in large communication systems (e.g. ad hoc wireless networks). An associated diffusion control problem is presented and we show that the value function of the N-particle controlled system converges to the value function of the limit diffusion control problem as $N \to \infty$. The diffusion coefficient in the limit model is typically degenerate, however under suitable conditions there is an equivalent formulation in terms of a controlled diffusion with a uniformly non-degenerate diffusion coefficient. Using this equivalence, we show that near optimal continuous feedback controls exist for the diffusion control problem, and then construct asymptotically optimal control policies for the N-particle systems based on such continuous feedback controls. Results for some preliminary numerical examples will be presented.

Poster proposal for the Stochastic Network Conference 2016 at UC San Diego

Balázs Gerencsér

February 4, 2016

Title of poster: Robust averaging - performance of the push-sum algorithm in the presence of transmission failures.

Name of presenter: Balázs Gerencsér (joint work with Julien Hendrickx)

Department and Institution of presenter: Probability & Statistics research division, Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences

Short abstract for the poster (300 words or less):

The push-sum algorithm allows distributed computing of the average of initial measurements on a network. It is designed to compute the perfect average on a directed graph, even when one is restricted to one-way and asynchronous communications.

We investigate the performance of this algorithm when the communication channels are unreliable meaning that messages can be lost. We show that convergence still holds, but the consensus value might be different from the true average.

We analyze the error of the final common value obtained. For the general case, we deduce fundamental properties that implicitly describe the distribution of the final value. Detailed theoretical and numerical discussion is presented for the essential case of two nodes.

The error performance is then compared with that of the standard consensus algorithm. This allows to see which one is the better choice to use based on the parameters of the actual network.

User equilibria in parallel Processor Sharing queues

Niffe Hermansson,
Department of Statistics,
University of Auckland

We study state based User Equilibria in systems of parallel queues in which each of the servers work according to the Egalitarian Processor Sharing service discipline. In this system a user is presented with several queues, each having the property that all users currently in the same queue are given simultaneous and identical service. The user now has to make a choice of which queue to join in order to minimize their own cost for passing through the system, without any consideration for the operation of the system in general. This choice is based on full knowledge of the system; the service rate of each server, the current occupancy of each queue, the arrival rate of other users as well as what choices the other users would make faced with each state of the system. In this context we are interested in the existence, properties and implications for system behaviour of Nash equilibrium policies. Previous work has indicated that even though this seems like a straight forward problem, it gives rise to counter intuitive results. This poster discusses the properties of user equilibria and the system behaviour under those equilibria while also exploring some of the perplexing properties of the system.

Title: Achievable Performance of Blind Policies in Heavy Traffic

Presenter: Bart Kamphorst
Institution: Centrum Wiskunde & Informatica (CWI), Amsterdam, the Netherlands
Department: Stochastics

Abstract: For a GI/GI/1 queue, we have shown that the average sojourn time under the (blind) Randomized Multilevel Feedback (RMLF) algorithm is no worse than that under the Shortest Remaining Processing Time algorithm times a logarithmic function of the system load ρ. Specifically, the gap in average sojourn time behaves like $O(\log(1/(1-\rho)))$ as ρ tends to one, and this behavior is tight for the M/M/1 queue. If the interarrival time and job size distributions are known, then there can be blind algorithms with a better performance; however, an appealing property of the RMLF algorithm is that its implementation does not depend on these input distributions.

The result is proven by combining techniques from two disparate areas: competitive analysis and applied probability. The proof only takes a first step into exploring the possibilities on the intersection of these areas; even so it is of independent interest.

Stable Transport Between Stationary Random Measures*

Ali Khezeli

March 27, 2016

Abstract

The problem of finding a transport kernel between (samples of) stationary random measures was initiated by Thorisson (1996)[1] with a necessary and sufficient condition for two random measures to be obtained by a random translation from each other, which is called a shift coupling. A particular example is when the second random measure is the Palm version of the first one. In the case of point processes a number of algorithms are presented to construct a shift coupling, which we will review. In particular, we are interested in the algorithm presented by Hoffman, Holroyd and Peres (2006)[2], which extends the Gale-Shapley stable marriage algorithm and the notion of stability to a continuum setting. We generalize this algorithm for arbitrary random measures, which is the work of my Ph.D. thesis. For this, we limit ourselves to *constrained* transport kernels, which is a special case of capacity constrained transport kernels. We give a definition of stability of constrained transport kernels and introduce a construction algorithm inspired by the Gale-Shapley stable marriage algorithm. For stable constrained transport kernels, we study existence, uniqueness, monotonicity w.r.t. the measures and boundedness.

*http://arxiv.org/abs/1504.02965

[1]Thorisson, H. (1996). Transforming random elements and shifting random fields. The Annals of Probability, 24(4), 2057-2064.

[2]Hoffman, C., Holroyd, A. E., & Peres, Y. (2006). A stable marriage of Poisson and Lebesgue. The Annals of Probability, 1241-1272

Title: Pathwise differentiability of reflected diffusions

Presenting author: David Lipshutz

Co-author: Kavita Ramanan

Affiliation: Division of Applied Mathematics, Brown University

Abstract: We consider pathwise differentiability of reflected diffusions in convex polyhedral domains, which arise as diffusion approximations in stochastic networks. We show that the directional derivatives can be characterized as reflected diffusions with time-varying domains and directions of reflection. As an application, we characterize derivatives of stochastic flows of reflected diffusions, and sensitivities of the reflected diffusions with respect to their drift and dispersion coefficients, and directions of reflection.

Synchronization of finite-state pulse-coupled oscillators

Hanbaek Lyu[1]

Department of Mathematics, The Ohio State University, Columbus, OH 43210 (yu.1242@osu.edu)

Abstract

We propose a novel generalized cellular automaton (GCA) model for discrete-time pulse-coupled inhibitory oscillators and study the emergence of synchrony. At each discrete time t, each vertex in a graph has a state in $\{0,\dots,\kappa-1\}$, and a special state $b(\kappa)=\lfloor\frac{\kappa-1}{2}\rfloor$ is designated as the 'blinking' state. At step t, simultaneously for all vertices, the state of a vertex increments from k to $k+1 \mod \kappa$ unless $k>b(\kappa)$ and at least one of its neighbors is in the state $b(\kappa)$. We call this model the κ-color *firefly cellular automaton* (FCA). On finite graphs, we obtain the following local global principle for tree networks: for $\kappa\in\{3,4,5,6\}$, the κ-color FCA on a tree synchronizes arbitrary initial configuration if and only if the maximum degree of the tree is less than κ [1].

For infinite graphs, we study FCA on $\mathbb{Z}^d$ starting from a random initial configuration drawn from the uniform product measure. FCA is closely related to other CA models for excitable media, namely, the cyclic cellular automaton (CCA) and the Greenberg-Hastings model (GHM). These models are well-studied on $\mathbb{Z}$ by considering their embedded particle system structure, which enables connections to random walks with i.i.d. increments. Along this line, it is known that CCA for $\kappa=3$ and GHM for any κ, the the probability of *not* having synchrony on any finite fixed interval at time t decays to zero in the order of $O(t^{-1/2+o(1)})$. We adopt similar methods and relate FCA on $\mathbb{Z}$ for any κ with locally dependent persistent random walks. We handle the local dependence in the increments via multiple methods and show that FCA belongs to the same universality class as above [3].

For $d\geq 2$, both CCA and GHM for any κ show the spontaneous emergence of self-organizing spirals in resemblance of the Belousov-Zhabotinsky reaction. These spirals divide the lattice into Voronoi-like cells, so the correlation length is finite and no local clustering occurs like in one dimension. Indeed, FCA for $\kappa\neq 4$ shows the similar behavior on $\mathbb{Z}^2$. However, this universal phenomenon of spiral formation is somehow inhibited in the four color FCA, and we do observe clustering on $\mathbb{Z}^2$. Even more surprisingly, this critical behavior does not seem to depend on the dimension d of the lattice. This may be thought of as a four-state Ising spin system with deterministic dynamics, with convergence toward global ferromagnetism.

As an application for clock synchronization in distributed networks, we generalize the four color FCA into a continuous model and show that on any finite tree with diameter d, arbitrary initial configuraiton $\Lambda_0: V\to S^1$ is synchronized in $O(d)$ times. Combining this model with a distributed spanning tree algorithm then gives a distributed algorithm for clock synchronization in a system of processors with distinct label and identical local clocks on any connected graph, which runs in $O(|V|)$ times [2].

This is a joint work with David Sivakoff.

References

[1] Lyu, H., 2015. Synchronization of finite-state pulse-coupled oscillators. Physica D: Nonlinear Phenomena 303, 28–38.
[2] Lyu, H., 2016. Phase synchronization of pulse-coupled excitable clocks. arXiv preprint arXiv:1604.08381.
[3] Lyu, H., Sivakoff, D., 2016. Synchronization of finite-state pulse-coupled oscillators on $\mathbb{Z}$. In preperation.

Poster details

- Title: Spatial processes with births and deaths - Hard-core regime.
- Name: Mayank Manjrekar.
- Department and Institution: Dept. of Mathematics, University of Texas at Austin.
- *Abstract:* We study Markov processes, where the state space is the set of discrete collection of points in $\mathbb{R}^n$. With time new points are born and existing points die due to interactions between the points. In particular, we consider Hard-core processes, where no two points can be within a distance 1 from each other. Whenever a point arrives and the hard-core condition is violated, some points are killed according to a transition kernel. We first show the existence of a process that locally evolves as described above. Further, if the points arrive according to a Poisson point process, existence and uniqueness of a stationary regime for the Markov process can be shown, for certain transition kernels. In the construction of the stationary regime we use a coupling from the past argument.

TITLE Universality of Power-of-d Load Balancing in Many-Server Systems

NAME Debankur Mukherjee

AFFILIATION Department of Mathematics and Computer Science,
Eindhoven University of Technology

ABSTRACT We consider a system of N parallel queues with unit exponential service rates and a single dispatcher where tasks arrive as a Poisson process of rate $\lambda(N)$. When a task arrives, the dispatcher assigns it to a server with the shortest queue among $d(N)$ randomly selected servers ($1 \leqslant d(N) \leqslant N$). This load balancing strategy is referred to as a JSQ($d(N)$) scheme, marking that it subsumes the celebrated Join-the-Shortest Queue (JSQ) policy as a crucial special case for $d(N) = N$. We develop a stochastic coupling construction to bound the difference in the queue length processes between the JSQ policy and a scheme with an arbitrary value of $d(N)$. We use the coupling to derive the fluid limit in case $d(N) \gg 1$ in the regime where $\lambda(N)/N \to \lambda < 1$ as $N \to \infty$, along with the associated fixed point. The fluid limit turns out not to depend on the exact growth rate of $d(N)$, and in particular coincides with that for the JSQ policy. We further leverage the coupling to establish that the diffusion limit in the critical regime where $(N - \lambda(N))/\sqrt{N} \to \beta > 0$ as $N \to \infty$ with $d(N) \gg \sqrt{N}\log(N)$ corresponds to that for the JSQ policy. These results indicate that the optimality of the JSQ policy can be preserved at the fluid-level and diffusion-level while reducing the overhead by nearly a factor $O(N)$ and $O(\sqrt{N})$, respectively.

POSTER APPLICATION

TITLE OF THE POSTER:

Reliability of energy networks under uncertainty: a large deviations approach

NAME OF THE PRESENTER:

Tommaso Nesti

INSTITUTION:

Centrum Wiskunde & Informatica (CWI), Stochastics Group

ABSTRACT:

The advent of renewable energy has huge implications for the design and control of power grids. With the increasing penetration of renewable sources, supply-side uncertainty is bound to grow dramatically. In this scenario, traditional reliability constraints have to be replaced by probabilistic guarantees.

An important constraint is that the temperature of each line should be bounded, and the typical manner in which this constraint is met is by imposing a certain bound on the current flowing in the line. However, since temperature responds gradually to current, from a rare events perspective the constraint on current is much more conservative than the constraint on temperature. Thus, imposing a constraint on the probability of temperature overload results in a larger line capacity, as compared to the same constraint on the probability of current overload.

In this poster we analyze the probability of said overloads and we develop capacity regions using large deviations techniques. We model stochastic power injections as solutions of stochastic differential equations involving a rarity parameter ϵ, which quantifies the amount of randomness in the system. Using the direct current (DC) approximation for the power flow equations, we apply the Freidlin-Wentzell theory to obtain sample path large deviations principles for the current and the temperature processes. This will provide estimates for the probability of overloads as $\epsilon \to 0$.

As the computation of the actual decay rate for the temperature process is a difficult infinite-dimensional varational problem, we focus on deriving approximations that are both fast-computable and improving on the traditional current constraint. We conclude by explicitely quantifying the capacity gain assuming Ornstein-Uhlenbeck processes for the power injections.

Title: Optimization of DNA sequencing using Stochastic Geometry

Elizabeth O'Reilly
Department of Mathematics, University of Texas at Austin

Abstract: The ability to quickly sequence DNA molecules has become indispensable for biological research. A widely used modern sequencing technology developed by Illumina uses the sequencing-by-synthesis method. DNA is fragmented into segments, which are then spread over a flow cell and amplified into clusters. Incorporation of complementary nucleotides results in a fluorescent signal that indicates the correct base pair. We model the configuration of fragments and cluster growth from these locations with a planar poisson point process and the Johnson-Mehl growth model. Noise in the signal results from internal phasing as well as interference from neighboring clusters, which we describe in terms of a shot-noise field. The yield of accurate reads is a function of the intensity of the point process and the radius of the clusters, which is determined by the amount of time they are allowed to grow. These design parameters are optimized to achieve a greater efficiency in the technology.

The Continuous Configuration Model: A Null for Community Detection on Weighted Networks

John Palowitch*

Statistics and Operations Research,
UNC Chapel Hill

Abstract

Community detection is the process of grouping strongly connected nodes in a network. Many community detection methods for unweighted networks have a theoretical basis in a null model, which provides an interpretation of resulting communities in terms of statistical significance. In this paper, we introduce a null for sparse weighted networks called the continuous configuration model. We prove a Central Limit Theorem for sums of edge weights under the model, and propose a community extraction method called CCME which combines this result with an iterative multiple testing framework. To benchmark the method, we provide a simulation framework that incorporates the continuous configuration model as a way to plant null or "background" nodes in weighted networks with communities. We show CCME to be competitive with existing methods in accurately identifying both disjoint and overlapping communities, while being particularly effective in ignoring background nodes when they exist. We present two real-world data sets with potential background nodes and analyze them with CCME, yielding results that correspond to known features of the data.

*JP has been partially supported by NIH R01 MH101819-01.

Detecting Markov Chain Instability: A Monte Carlo Approach

Brendan Patch*†‡

We devise a Monte Carlo based method for detecting whether a non-negative Markov chain is stable for a given set of potential parameterizations. More precisely, for a given set in parameter space, we develop an algorithm that is capable of deciding whether the set has a subset of positive Lebesgue measure for which the Markov chain is unstable. The approach is based on a variant of simulated annealing, and consequently only mild assumptions are needed to obtain performance guarantees.

This is joint work with Michel Mandjes (University of Amsterdam) and Neil Walton (The University of Manchester).

*ARC Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS), The School of Mathematics and Physics, The University of Queensland, Qld 4072, Australia.

†Korteweg-de Vries Institute for Mathematics, University of Amsterdam

‡Email: b.patch@uq.edu.au

Jamol Pender Poster Submission

Title: Strong Approximations for Time Varying Queues with Non-Renewal Arrival and Service Processes

Abstract: In real stochastic systems, the arrival and service processes may not be renewal processes. For example, in many telecommunication systems such as internet traffic where data traffic is bursty, the sequence of inter-arrival times and service times are often correlated and dependent. One way to model this non-renewal behavior is to use Markovian Arrival Processes (MAP's) and Markovian Service Processes (MSP's). MAP's and MSP's allow for inter-arrival and service times to be dependent, while providing the analytical tractability of simple Markov processes. To this end, we construct a Poisson process representation for the MAP/PH/infinity queueing process dynamics and leverage strong approximations for Poisson processes. As a result, the fluid and diffusion limit theorems illuminate how the dependence structure of the arrival or service processes can affect the sample path behavior of the queueing process. Finally, our Poisson representation for MAP's is useful for simulation purposes and may be of independent interest.

Moment-Based Spectral Analysis of Random Graphs with Given Expected Degrees

Presenter: M. Amin Rahimian
Electrical and Systems Engineering
University of Pennsylvania.

Keywords: Complex Networks, Random Graph Models, Spectral Graph Theory, Random Matrix Theory

Abstract: In this paper, we analyze the limiting spectral distribution of the adjacency matrix of a random graph ensemble, proposed by Chung and Lu, in which a given expected degree sequence $\overline{w}_n^{\mathsf{T}} = (w_1^{(n)}, \ldots, w_n^{(n)})$ is prescribed on the ensemble. Let $a_{i,j} = 1$ if there is an edge between the nodes $\{i, j\}$ and zero otherwise, and consider the normalized random adjacency matrix of the graph ensemble: $\boldsymbol{A}_n = [a_{i,j}/\sqrt{n}]_{i,j=1}^n$. The empirical spectral distribution of $\boldsymbol{A}_n$ denoted by $\mathbf{F}_n(\cdot)$ is the empirical measure putting a mass $1/n$ at each of the n real eigenvalues of the symmetric matrix $\boldsymbol{A}_n$. Under some technical conditions on the expected degree sequence, we show that with probability one, $\mathbf{F}_n(\cdot)$ converges weakly to a deterministic distribution $F(\cdot)$. Furthermore, we fully characterize this distribution by providing explicit expressions for the moments of $F(\cdot)$. We apply our results to well known degree distributions, such as power-law and exponential. The asymptotic expressions of the spectral moments in each case provide significant insights about the bulk behavior of the eigenvalue spectrum.

Title of Poster:
A Stochastic Programming Based Approach to Control of Assemble-to-Order Inventory Systems

Presenter:
Martin I. Reiman

Department and Institution of presenter:
Industrial Engineering and Operations Research, Columbia University

Abstract:

The assemble-to-order (ATO) system is a classical model in inventory theory, where multiple components are used to produce multiple products. All components are obtained from an uncapacitated supplier after a (component dependent) deterministic lead time, while demand for the products forms a compound Poisson process. Assembly is assumed to be performed instantaneously, so all inventory is held as components rather than finished products. Demand that is not met immediately is backlogged. The optimal control for this system (where the goal is to minimize the long run average inventory + backlog cost) is not known except for some very special cases.

This poster presents an approach to this problem that is based on the solution of a related stochastic program (SP). When all lead times are identical this SP is a relaxation of a two stage SP that was previously put forward as a one-period model of an ATO system. More generally, when there are K distinct lead times, the SP has K+1 stages. We have shown that the SP provides a lower bound on the achievable cost under any feasible policy. We have also provided a ⊢translation ⊣ of the solution of the SP into a policy for the inventory system. When all components have the same lead time we have shown that this policy is asymptotically optimal as the lead time grows large. We also have shown asymptotic optimality in the non-identical lead time case when the differences of the lead times are small compared to the lead times themselves. Implementing our policy requires solving the SP. Along these lines we have obtained easily computable solutions for some special cases, and have shown numerically, for these cases, that our policy performs well. We have also obtained some structural results that make solving the SP simpler under certain conditions.

(Based on joint work with Mustafa Dogru, Qiong Wang and Haohua Wan)

Sensitivity analysis for Markov chains

Chang-Han Rhee

(Joint work with Peter Glynn)

Centrum Wiskunde & Informatica

June 20, 2016

Abstract

Let $X(\theta) = (X_n(\theta) : n \geq 0)$ be a parametrized family of Markov chains. We introduce a general theory that provides sufficient conditions for differentiability (w.r.t. θ) of various types of performance measures such as the ones associated with discounted random horizon expectations $\mathbb{E}\sum_{j=0}^{\tau} e^{\sum_{i=0}^{j} g(X_i(\theta))} f(X_j(\theta))$ and stationary expectations $\mathbb{E}f(X_\infty(\theta))$, which are known to be notoriously difficult to deal with. Such derivatives are valuable in various contexts such as understanding the relations between input and output variables in a system, testing the robustness of a model with respect to parameter uncertainty, and model simplification. They also play a key role when one is numerically optimizaing the performance measures over the decision parameter θ. The new theory is easy to apply on the basis of the model building blocks and requires weaker conditions than the ones provided by the existing literature. We also identify expressions for the derivatives that are useful for computation via simulation.

Stochastic network thinking applied to firing patterns of stellate neurons

Peter Rowat
Institute for Neural Computation,
University of California San Diego,
La Jolla, CA 92093-0523, USA

Priscilla Greenwood
Mathematics Department
University of British Columbia
Vancouver, BC, Canada V6T 1W5

The medial entorhinal cortex (MEC) provides an important input to the hippocampus, which plays a major role in memory and spatial navigation. The primary excitatory cells of the MEC are stellate neural cells which feed forward to place cells in the hippocampus. Most stellate cells are grid cells, which play an important role in spatial navigation in rats, as well as in bats and humans. Curiously, grid cells fire when the rat traverses points in space that form an hexagonal pattern, a response that must be learned by a stochastic network of cells. It is established that the ionic current known as the "I_h" current plays an important role in the dynamics of grid cells. Using a reduced dynamic model of stellate cells we explore the role that the I_h current plays in this phenomenon.

Poster for Stochastic Networks Conference 2016

Title of Poster: Optimality of the Generalized $c\mu$ Rule
in the Moderate Deviation Regime.

Name of Presenter: Subhamay saha

Affiliation: Postdoctoral Fellow,
Department of Electrical Engineering,
Technion - Israel Institute of Technology, Haifa, Israel.

Abstract: This paper studies a multiclass queueing system with an associated risk-sensitive cost observed in heavy traffic at the moderate deviation scale, accounting for convex queue length penalties. The main result is the asymptotic optimality of a dynamic index policy known from the diffusion scale heavy traffic literature as the generalized $c\mu$ rule.

Application poster session

Title of poster:	Asymptotic analysis of a cascading failure model
Name of presenter:	Fiona Sloothaak
Institution:	Eindhoven University of Technology
Department:	Mathematics and Computer Science

Abstract

Large blackouts of electric power transmission systems have major impact on modern-day society. In order to analyze the reliability of energy networks, we use a probabilistic approach to study the occurrence of severe blackouts caused by a single failure. For a stylized cascading failure model, we assume that an electricity transmission line trips if its capacity is insufficient for the load surge induced by the single failure. Every tripped line again causes additional loading on the remaining lines, resulting in a cascading failure effect. The risk of cascading effects turns more severe when loading increases, and at criticality, the blackout size has power-law behavior. Analysis of the stylized model is possible by utilizing results from extreme value theory, deriving asymptotics for analytical expressions, and relating theory of exit times from moving boundaries for stable random walks.

Title of the Poster:

Convergence Rates in Distributed Learning: Acceleration, Network Independence and Uniform Social Sampling.

Name of the Presenter:

Cesar A Uribe

Affiliation:

Coordinated Science Laboratory – University of Illinois at Urbana-Champaign

Abstract:

Networks in which individuals share and aggregate information following local rules and communication constraints are known as 'social networks.' The objective is to achieve global behaviors by repeatedly aggregating local information without complete knowledge of the network. Although the name is reminiscent of human behavior, a network might well be made of sensors, robots or other engineered system. Although many results on these themes have appeared in recent years, the study of distributed decision making and computation can be traced back to the classic papers from the 70s and 80s. In this work we consider the problem of distributed learning where a network of agents repeatedly observe some random processes and would like to collectively agree on a hypothesis that best explains all the observations in the network. In addition agents also repeatedly receive information from their neighbors, defined over a time-varying sequence of graphs.

We present a sequence of distributed learning protocols and show they allow the network to learn the set of hypotheses which explain the data collected by all the nodes best (i.e. consistency). Moreover we provide geometric, non-asymptotic, and explicit characterization of their convergence rates, which immediately leads to finite-time bounds that scale intelligibly with the number of nodes. We consider specific protocols for static-undirected and time-varying (un)directed graphs. In the case of directed graphs we show that a push-sum inspired algorithm allows for network independent rates. For fixed undirected graphs we show the proposed protocol converges a factor of n faster than previous algorithms. Finally, we show protocols with uniform social sampling reduce the amount of communication needed at each time step but reduces the converge rate in a factor inversely proportional to the number of hypothesis.

Selfish routing in a network of parallel queues

Yijun(Alex) Wang*

Department of Statistics,

The University of Auckland

Abstract. Consider a network where two types of routes are available for users to choose from a source to a destination. On one type of route (FIFO), delay increases as traffic increases and on the other (batch-service), the delay decreases as traffic increases. The user chooses their own route through the network, in order to minimize the expected delay. Adding capacity to queues does not always yield better performance, which is known as Downs-Thomson paradox. This poster discusses the existence, uniqueness and properties of user equilibrium under both probabilistic and state-dependent routing with examples.

Keywords: Queueing network · User equilibria · Parallel queues · Downs-Thomson paradox · Wardrop equilibrium

*Department of Statistics, The University of Auckland, Private Bag 92019, Auckland, New Zealand

Stochastic Networks Conference 2016
Poster Session Application

1 Title of poster

Adaptive Policies for Scheduling with Reconfiguration Delay: An End-to-End Solution for All-Optical Data Centers

2 Name of presenter

Chang-Heng Wang

3 Department and Institution

University of California, San Diego
Department of Electrical and Computer Engineering

4 Abstract

All-optical switching networks have been considered a promising candidate for the next generation data center networks thanks to its scalability in data bandwidth and power efficiency. However, the bufferless nature and the nonzero reconfiguration delay of optical switches remain great challenges in deploying all-optical networks. This work considers the end-to-end scheduling for all-optical data center networks with no in-network buffer and nonzero reconfiguration delay. A framework is proposed to deal with the nonzero reconfiguration delay. The proposed approach constructs an adaptive variant of any given scheduling policy. It is shown that if a scheduling policy guarantees its schedules to have schedule weights close to the MaxWeight schedule (and thus is throughput optimal in the zero reconfiguration regime), then the throughput optimality is inherited by its adaptive variant (in any nonzero reconfiguration delay regime). As a corollary, a class of adaptive variants of the well known MaxWeight policy is shown to achieve throughput optimality without prior knowledge of the traffic load. Furthermore, through numerical simulations, the simplest such policy, namely the Adaptive MaxWeight (AMW), is shown to exhibit better delay performance than prior work.

Title: **Optimal Routing to Remote Queues**

Presenter: Yao Yu (于耀)

Affiliations: Department of Industrial and Systems Engineering, NC State University

Abstract

We develop optimal routing policies for remote queueing systems, in which each arrival, after being routed to join one of several single-server queues in parallel, will experience a *pre-arrival delay*. Motivated by service systems in which system state (e.g., queue length and waiting time) is available for routing decisions, we intend to use pre-arrival delays to model commute times of arrivals, such as patients' transportation times before arriving at clinics and data packets' transmission times to web servers. For parallel queues with no pre-arrival delays, it is well known that the *join-the-shortest queue* (JSQ) routing policy is asymptotically optimal in minimizing the total queue length and in-queue waiting time. In addition, under JSQ the performance of the parallel system with no pre-arrival delays is asymptotically equivalent to the pooled system in heavy traffic. In the presence of pre-arrival delays, unfortunately, JSQ can be disastrous for the system performance, causing excessively large in-queue waiting time and an undesired bouncing effect (having negatively correlated queue performance). In order to reduce the waiting time and to minimize the total queue length, we propose a new state-dependent probabilistic routing policy, named *JSQ with a root-excess bias* (JSQ-REB). Specifically, JSQ-REB means that we route a customer to the shortest queue with a slight bias that is proportional to the square root of the fraction of idleness (i.e., one minus the traffic intensity). We prove a heavy-traffic limit theorem by showing that, under the proposed JSQ-REB policy, the parallel system is asymptotically equivalent to the pooled system as the traffic intensity approaches one, so that both the waiting time and total queue length are minimized, and the excessive congestion caused by pre-arrival delays is eliminated.

Joint work with Shuangchi He from NUS and Yunan Liu from NCSU.

Title: Minimizing heat dissipation in DC networks using batteries

Presenter: Alessandro Zocca

Affiliation: CWI Amsterdam

Abstract: Energy transmission networks dissipate a non-negligible fraction of the power they transport due to the heat loss in the transmission lines. In this work we explore how the transport of energy can be made more efficient by adding to the network multiple batteries that can coordinate their operations. Such batteries can both charge in case of excess of power generation in the network or discharge to meet the network demand. Either way, the presence of batteries in the network can be leveraged to mitigate the intrinsic uncertainty in the power consumption and generation and, hence, transport the energy more efficiently through the network.
We consider a resistive DC network with stochastic power injections and show how the total average heat dissipation depends on the network structure and on the batteries operations. Furthermore, in the case where the power injections are modeled by Ornstein-Uhlenbeck processes, we derive the dynamical optimal control for the batteries over a finite time interval.

Title of poster: Delay, memory and messaging tradeoffs in distributed service systems.

Name of presenter: Martin Zubeldia.

Department and institution of presenter: Laboratory for Information and Decision Systems (LIDS), Department of Electrical Engineering and Computer Science (EECS), Massachusetts Institute of Technology (MIT).

Short abstract: We consider the following distributed service model: jobs with unit mean, exponentially distributed and independent processing times arrive as a Poisson process of rate λN, with $0 < \lambda < 1$, and are immediately dispatched to a queue associated with one of N identical servers. We assume that the dispatching decisions are made by a central dispatcher endowed with a finite memory, and with the capacity to exchange messages with the servers.

We study the fundamental limits in the amount of resources (memory bits and message exchange rate) needed, in order to drive the expected steady state queueing delay of a typical job to zero, as N increases. We construct a policy which drives the delay to zero when either (i) the message rate grows superlinearly with N, or (ii) the memory grows superlogarithmically with N. The analysis of this policy uses the fluid model technique. Moreover, we show that any policy that has a certain symmetry property, and for which neither condition (i) or (ii) holds, results in an expected queueing delay which is bounded away from zero.

Finally, using the fluid limit approach once more, we show that when our policy only utilizes a linear message rate (αN) and a logarithmic memory, the limiting expected queueing delay (although positive) is uniformly upper bounded as $\lambda \uparrow 1$, for any linear message rate (i.e., for any $\alpha > 0$), no matter how small. This is a significant improvement over the popular "power-of-d-choices" policy, which has a limiting expected delay that grows as $\log\left(\frac{1}{1-\lambda}\right)$ when $\lambda \uparrow 1$.