Stochastic Networks Conference
Titles and Abstracts of Plenary Talks

UC San Diego

June 20-24, 2016

1 MONDAY, JUNE 20, 2016

1.1 MONDAY MORNING

Session Chair: Cristina Costantini (Universita di Chieti-Pescara)

9.45-10.45 a.m.

Tom Kurtz **(U. Wisconsin)**

Exchangeable systems and mean field approximations for stochastic networks

Drawing examples from queueing and neural network models, general approaches for deriving mean field approximations for complex models will be described. Central to this approach is the observation that if a sequence of exchangeable families (finite or infinite) converges in distribution to an infinite exchangeable family, then the corresponding sequence of de Finetti measures converges to the de Finetti measure of the limit. If the random variables are cadlag stochastic processes, then under minor restrictions, the measure valued stochastic process given by the de Finetti measure of the process states at each time t converges in distribution to the measure-valued process corresponding to the limiting infinite system. In the Markov setting, uniqueness for the martingale problem corresponding to the limiting infinite system implies uniqueness for the martingale problem (appropriately defined) for the limiting measure-valued process.

11.15 a.m.-12.15 p.m.

Philippe Robert **(INRIA)**

Asymptotics of stochastic protein assembly models

Self-assembly of proteins is a biological phenomenon which gives rise to spontaneous formation of amyloid fibrils or polymers. A striking feature of this process is that the starting point of the chemical reaction, called nucleation, exhibits an important variability among replicated experiments. In the case of neuro-degenerative diseases, such as Alzheimer or Prion diseases, this variation can be of the order of several years. Two populations of chemical components are involved: monomers (proteins) and polymers. Initially there are only monomers. There are two reactions for the polymerization of a monomer: either two monomers collide to combine into a polymer or a monomer is polymerised after the encounter of a polymer. After a general introduction, the talk will present several stochastic models of increasing complexity from a basic two-dimensional Markovian model to more sophisticated multi-dimensional processes. The problem of fitting parameters will be presented with data from experiments in biochemistry. Mathematically, the techniques use classical stochastic calculus methods, stochastic averaging principles and a set of classical ODE's, the Becker-Doring equations. It is shown that a set of different time scales play an important role for these models.

The talk is based on joint works with Marie Doumic, Sarah Eugène, Wen Sun and Wei-Feng Xue.

1.2 MONDAY AFTERNOON

2-3 p.m. *Chair:* Tara Javidi (UCSD)

Fan Chung Graham **(UC San Diego)**
Graph coloring games and voter models

We investigate certain random processes on graphs which are related to the so-called Tsetlin library random walks as well as to some variants of the classical voter model. A specific example can be described as a hypergraph coloring game. Each node represents a voter and is colored according to its preferred candidate, or undecided. Each hyperedge is a subset of nodes and can be viewed as a chat group. In each round of the game, one chat group is chosen randomly, and voters in the group can change colors based on interactions. We analyze the game as a random walk on the associated weighted directed state graph. Under certain 'memoryless' conditions, the spectrum of the state graph can be explicitly determined by using semigroup spectral theory. It can be shown that random walk on the state graph converges to its stationary distribution in $O(m \log n)$ time, where n denotes the number of voters and m denotes the number of chat groups. This can then be used to determine the appropriate cut-off time for voting. We also consider a partial memoryless version which can be used to approximate general voter games.

This talk is based on joint papers with Ron Graham.

3.30-7.30 p.m.
POSTER PRESENTATIONS AND POSTER SESSION
Chair: Amber Puha (CSUSM)

See separate listing of poster titles and abstracts.

2 TUESDAY, JUNE 21, 2016

2.1 TUESDAY MORNING

Session Chair: Adam Wierman (Cal Tech)

9.30-10.30 a.m.
Laurent Massoulie **(Microsoft Research-INRIA)**
Non-regular Ramanujan graphs and community detection

This talk will review the notion of Ramanujan graphs and the result of Alon and Boppana according to which such graphs have a maximal spectral separation. It will then present a spectral characterization of non-backtracking matrices of random graphs sampled according to the so-called stochastic block model. This characterization implies that i) Erdos-Rényi graphs satisfy a property extending the notion of Ramanujan graph to the non-regular case, and ii) spectral clustering based on the non-backtracking matrix can perform non-trivial community detection down to the so-called Kesten-Stigum threshold. This last property establishes the "spectral redemption" conjecture made by Krzakala et al.

This talk is based on joint work with Charles Bordenave and Marc Lelarge

11 a.m.-noon
Emmanuel Abbe **(Princeton)**
Reaching and crossing thresholds in the stochastic block model

This talk overviews our recent results on the stochastic block model thresholds, both for exact recovery (the CH-divergence threshold) and weak recovery (the KS threshold). In particular, it is shown how a linearized acyclic BP algorithm - or equivalently a spectral method on a nonbacktracking operator of generalized order - achieves the KS threshold in quasi-linear time, while from 5 communities, it is possible to cross the KS threshold with information theory.

This talk is based on joint works with C.Sandon:
`http://arxiv.org/abs/1503.00609` and `http://arxiv.org/abs/1512.09080`.

2.2   TUESDAY AFTERNOON

Session Chair: Amarjit Budhiraja (Univ. North Carolina, Chapel Hill)

2-3 p.m.

Bert Zwart **(Centrum Wiskunde & Informatica)**

Sample-path large deviations for heavy tails

For rare events in statistical mechanics and related topics, a beautiful theory of large deviations exists. However, many rare events in man-made networks exhibit heavy-tailed features. Examples are file sizes, delays and financial losses, but also magnitudes of systemic events, such as the size of a blackout in a power grid. The theory developed in the heavy-tailed case is mostly restricted to model-specific results as well as results pertaining to events that are caused by a single big jump. In this talk we develop sample-path large deviations for random walks and Levy processes in the heavy-tailed case that go beyond such restrictions.

Joint work with Jose Blanchet (Columbia Univ) and Chang-Han Rhee (CWI).

3.30-4.30 p.m.

Mark Lewis **(Cornell)**

Dynamic control of a tandem queueing system with abandonments

We consider an extension of classic scheduling problems for a tandem queueing system by including customer impatience. In such scenarios, the server(s) must balance abandonments from each phase of service with the need to prioritize higher reward customers. This presents an interesting challenge since the trade-off between the cost of abandonments and revenue maximization are not at all clear.

The model closely replicates the decisions made by a medical service provider in the new "triage-treatment-and-release" system begun in Lutheran Medical Center in Brooklyn, NY. In order to mitigate the concern that waiting times are increased by prioritizing one phase or the other, we introduce K-level heuristics that balance the two. Numerical examples show the efficacy of these heuristics.

3 WEDNESDAY, JUNE 22, 2016

3.1 WEDNESDAY MORNING

Session Chair: Andres Ferragut (Universidad ORT Uruguay)

9.30-10.30 a.m.
Angelia Nedich **(U. Illinois, Urbana-Champaign)**
Decentralized hypothesis testing on graphs

We will consider the problem of distributed cooperative learning in a network of agents, where the agents are repeatedly gaining partial information about an unknown random variable whose distribution is to be jointly estimated. The learning is based on Bayesian update adapted to distributed information structure inherited from the network. The joint objective of the agent system is to globally agree on a hypothesis (distribution) that best describes the observed data by all agents in the network. Interactions between agents occur according to an unknown sequence of time-varying graphs. We highlight some interesting aspects of Bayesian learning and stochastic approximation approach for the case of a single agent, which has not been observed before and it allows for a new connection between optimization and statistical learning. Then, we discuss and analyze the general case where subsets of agents have conflicting hypothesis models, in the sense that the optimal solutions are different if the subset of agents were isolated. Additionally, we provide a new non-Bayesian learning protocol that converges an order of magnitude faster than the learning protocols currently available in the literature for arbitrary fixed undirected graphs. Our results establish consistency and a non-asymptotic, explicit, geometric convergence rate for the learning dynamics.

11 a.m.-noon
Guy Bresler **(MIT)**
Learning a tree-structured Ising model in order to make predictions

We study the problem of learning a high-dimensional tree graphical model from samples such that low-order marginals are accurate. This enables making accurate predictions based on partial observations. A take-home message is that the necessary number of samples is dramatically lower than is required to learn the exact underlying tree. On the technical side, this requires reasoning about predictions computed using a completely incorrect tree. Joint work with Mina Karzand.

3.2 WEDNESDAY AFTERNOON

2-3 p.m. *Chair:* Masakiyo Miyazawa (Tokyo University of Science)
Maury Bramson **(U. of Minnesota)**
Stability and (mostly) instability of the MaxWeight policy

Consider a switched queueing network with general routing among its queues. The MaxWeight policy is the service rule that assigns available service by maximizing the objective function $\sum Q_j \sigma_j$ among the different feasible service options, where Q_j denotes queue size and σ_j denotes the amount of service to be executed at that queue. MaxWeight is a greedy policy that does not depend on knowledge of arrival rates and is not difficult to implement. These properties, as well as its simple formulation, suggest MaxWeight as a serious candidate for implementation in the setting of switched queueing networks.

It turns out to be the case, however, that MaxWeight is not maximally stable in general, even in the single class setting. A fluid model variant of MaxWeight was shown by Andrews-Zhang (2003) not to be maximally stable. Here, we discuss work in progress with B. DAuria and N. Walton that shows MaxWeight itself is not maximally stable in a variety of settings. We also discuss results that show MaxWeight is maximally stable in a much more restrictive setting.

3.30-4.30 p.m. *Chair:* R. Srikant (UIUC)
Yi Lu **(U. Illinois, Urbana-Champaign)**
Scheduling in Hadoop clusters: Affinity scheduling with delay optimality but no planning

We found that scheduling in Hadoop clusters belongs to the class of affinity scheduling problems, where processing rates depend on the task-server pair. Delay-optimal algorithms in heavy-traffic were first proposed by Harrison-Lopez (1999) and Bell-Williams (2005). However, the algorithms require the arrival rates of each task type, which are difficult to obtain in data centers. On the other hand, Stolyar (2004) proposed to use the MaxWeight algorithm, which does not require arrival rates, but also loses the delay-optimal property, that is, MaxWeight does not minimize linear functions of delay.

We propose an algorithm that does not require arrival rates, but still achieves delay-optimality in the Hadoop cluster setting. The proof is based on the construction of an ideal load decomposition. We will also discuss connections to prior work. The algorithm is implemented in a Hadoop scheduler and achieves > 10 times improvement over existing schedulers. This talk is based on joint works with Q. Xie and A. Yekkehkhany.

See next page for one more talk on Wednesday afternoon.

4.30-5.30 p.m. (Wednesday) *Chair:* R. Srikant (UIUC)
Siva Maguluri **(IBM T. J. Watson Research Center)**
Delay-optimal scheduling for data center networks and input-queued switches in heavy traffic

Today's era of cloud computing is powered by massive data centers. A data center network enables exchange of data in the form of packets among the servers within these data centers. Given the size of today's data centers, it is desirable to design low complexity scheduling algorithms which result in a fixed average packet delay, independent of the size of the data center. We consider the scheduling problem in an input-queued switch, which is a good abstraction for a data center network. In particular, we study the queue length (equivalently, delay) behavior under the so-called MaxWeight scheduling algorithm, which has low computational complexity. Under various traffic patterns, we show that the algorithm achieves optimal scaling of the heavy-traffic scaled queue length with respect to the size of the switch. This settles one version of an open conjecture that has been a central question in the area of stochastic networks. We obtain this result by using a Lyapunov-type drift technique to characterize the heavy-traffic behavior of the expected total queue length in the network, in steady-state.

4 THURSDAY, JUNE 23, 2016

4.1 THURSDAY MORNING

Session Chair: Marty Reiman (Columbia U.)

9.30-10.30 a.m.
Kavita Ramanan **(Brown U.)**
Ergodicity of stochastic networks

Common themes in the study of stochastic networks include the study of ergodicity of stochastic processes that describe the underlying dynamics, and characterization of the limit of stationary distributions in a relevant asymptotic regime. We will consider several problems arising in stochastic networks that involve either an infinite-dimensional representation or degenerate dynamics, where more standard techniques to establish ergodicity such as Harris recurrence do not easily apply. We will describe alternative approaches that have proved useful in these situations, and illustrate their application to several concrete examples.

11 a.m.-noon
Rami Atar **(Technion)**
On the measure-valued Skorohod map

A Skorohod-type transformation that acts on paths with values in the space of measures over the real line is argued to provide a generic model for priority. We will describe how it applies to the single-server earliest-deadline-first and shortest-job-first, and discuss how its generalizations might be useful for other models.

This talk is based on joint work with Biswas, Kaspi and Ramanan.

4.2   THURSDAY AFTERNOON

1.30-2.30 p.m. *Chair:* Ilze Ziedins (U. Auckland)
Galit Yom Tov **(Technion)**
The impact of delay announcements on hospital network coordination and waiting times

We investigate the impact of delay announcements on the coordination within hospital networks using a combination of empirical observations and numerical experiments. We show that patients take delay information into account when choosing emergency service providers and that such information can help increase coordination in the network, leading to improvements in performance of the network, as measured by Emergency Department wait times. Our numerical results indicate that the level of coordination that can be achieved is limited by the patients sensitivity to waiting, the load of the system, the heterogeneity among hospitals, and, importantly, the method hospital use to estimate delays. We show that delay estimators that are based on historical average may cause oscillation in the system and lead to higher average waiting times when patients are sensitive to delay. We provide empirical evidence which suggests that such oscillations occur in hospital networks in the US.

2.30-3.30 p.m. *Chair:* Ilze Ziedins (U. Auckland)
Neil Walton **(Manchester University)**
Join-the-shortest-queue and idle-one-first in the NDS regime

We discuss recent and ongoing work on parallel queueing models. Here the Join-the-Shortest-Queue policy is known to be optimal. However, in the Halfin-Whitt regime, the many-queue limit of this queueing system is degenerate. Thus, in order to obtain a more meaningful asymptotic description, we instead consider a non-degenerate slow-down regime. We characterize the diffusion limit obtained and analyse it to construct a simple, low message rate policy called Idle-One-First. This policy is optimal in the NDS regime. Further, we discuss associated insensitivity properties found for JSQ in the NDS limit.

The presentation is joint work with Varun Gupta, University of Chicago.

See next page for one more talk on Thursday afternoon.

4-5 p.m. (Thursday) *Chair:* Amy Ward (USC)
Itai Gurvich (Northwestern)
Beyond heavy-traffic regimes: Universal bounds and controls for the single-server queue

Brownian approximations are widely studied because of their tractability relative to the original queueing models. Their stationary distributions are used as proxies for those of the original queues that they model and the convergence of suitably scaled and centered processes provides mathematical support for the use of these Brownian models. To establish convergence, one must impose assumptions directly on the primitives or, indirectly, on the parameters of a related optimization problem. These assumptions reflect an interpretation of the underlying parameters—a classification into so-called heavy-traffic regimes that specify a scaling relationship between the utilization and the arrival rate. From a heuristic point of view, though, there is an almost immediate Brownian analogue of the queueing model that is derived directly from the primitives and requires no (limit) interpretation of the parameters. In this paper we prove that for the fundamental $M/GI/1+GI$ queue, the direct intuitive (limitless approach) in fact works. The Brownian model is universally (i.e., across regimes and patience distributions) accurate. It maintains the tractability and appeal of the limit approximations while avoiding many of the assumptions that facilitate them. In the process of building mathematical support for the accuracy of this model, we introduce a framework built around "queue families" that allows us to treat various patience distributions simultaneously, and uncovers the role of a concentration property of the queue.

5.30-8.30 p.m.
CONFERENCE DINNER, UC San Diego Faculty Club.
Dinner tickets must be purchased in advance, with registration.

5 FRIDAY, JUNE 24, 2016

5.1 FRIDAY MORNING

Session Chair: John Hasenbein (U. Texas, Austin)

9.30-10.30 a.m.

Francois Baccelli **(U. Texas, Austin)**

Navigation on a stochastic network (For a related paper, click here).

The talk will be focused on dynamics on discrete random structures. We will first consider compatible dynamics on the points of a stationary point process, namely "rules to navigate from point to point" which are preserved by translations. Each such dynamics defines a random graph on the points of the process. The connected components of this graph can be split into a collection of foils, which are the analogue of the stable manifold of the dynamics. We will give a general classification of such dynamics in terms of the cardinality of the foils of these connected components. There are three types: F/F (finitely many finite foils), I/F (infinitely many finite foils), and I/I. We will then consider compatible dynamics on random graphs and random networks (marked graphs), namely "rules to navigate from node to node" which are preserved by graph or network isomorphisms. We will show that this classification also holds for all unimodular random graphs and networks. We will complement this by analytical results on the relative intensities of the foils. These results will be illustrated by concrete examples of navigation rules, both on point processes and on random networks.

Joint work with M.-O. Haji-Mirsadeghi, Department of Mathematics, Sharif University, and A. Khezeli, Department of Mathematics, UT Austin.

11 a.m.-noon

Jim Dai **(Cornell)**

Diffusion approximations of queueing networks: a direct method for proving stationary distribution convergence

Diffusion processes such as semimartingale reflecting Brownian motions (SRBM) have been widely used to approximate performance processes of queueing networks in heavy traffic. In this talk, I will promote a direct method, first introduced by Miyazawa, for proving the convergence of the stationary distributions of these performance processes. The direct method starts with the basic adjoint relationships for the pre-limit stationary distributions, and the tightness of these distributions is verified through a system of linear inequalities. I will show the method's effectiveness in the generalized Jackson network setting and discuss its potential for multiclass queueing networks. The direct method contrasts sharply with the limit interchange method, an indirect method pioneered by Gamarnik and Zeevi (2006) and later used by Gurvich (2014), Ye and Yao (2016), and many others.

This talk is based on joint works with Anton Braverman and Masakiyo Miyazawa.