Conference on Stochastic Processing Networks
in Honor of J. Michael Harrison
Graduate School of Business, Stanford University
August 29-30, 2009

ABSTRACTS

Speaker: Jim Dai
Title: Maximum Pressure Policies in Stochastic Processing Networks
Abstract: Stochastic processing networks have been introduced in a series of three papers by Harrison (2000, 2002, 2003). These networks are much more general than multi-class queueing networks. The added generality allows one to model skill-based routing in call centers, operator-machine interactions in semiconductor wafer fabrication facilities, and combined input- and output-queued data switches.

We introduce a family of maximum pressure policies that can be used to operate a stochastic processing network. Such a policy is shown to be throughput optimal for a wide class of networks. When the holding cost rate is a quadratic function of the buffer contents, we show that a maximum pressure policy asymptotically minimizes the holding cost under a heavy traffic condition and a complete resource pooling condition.

[This talk is based on the joint work with Wuqin Lin at Northwestern University.]

Speaker: Maury Bramson, University of Minnesota
Title: Positive Recurrence of Reflecting Brownian Motion in 3 Dimensions
Abstract: Precise conditions are known for positive recurrence of semimartingale reflecting Brown motion (SRBM) in 2 dimensions. The behavior in 4 and more dimensions is largely unexplored. In 3 dimensions, sufficient conditions for positive recurrence are given by El Kharroubi et al. (2002). In joint work with Jim Dai and Mike Harrison, it was recently shown that these results are, in fact, sufficient. In this talk, we summarize the above material.

Speaker: Steve Shreve, Carnegie Mellon University
Title: Double Skorokhod Map and Reneging Real-Time Queues
Abstract: An explicit formula for the Skorokhod map on [0,a] is provided. Specifically, it is shown that on the space of right-continuous functions with left limits taking values in the real numbers,

$$\psi(t) - \left[(\psi(0)-a)^{+} \wedge \inf_{u\in[0,t]} \psi(u)\right] \vee \sup_{s\in[0,t]} \left[(\psi(s)-a) \wedge \inf_{u\in[s,t]} \psi(u)\right].$$

is the unique function taking values in [0,a] that is obtained from ψ by minimal “pushing” at the endpoints 0 and a. An application of this result to real-time queues with reneging is outlined.

[This is joint work with L. Kruk, J. Lehoczky and K. Ramanan.]

Speaker: Marty Reiman, Bell Laboratories
Title: Newsvendor Networks and Assemble-to-Order Inventory Systems
Abstract: In the assemble-to-order (ATO) inventory system multiple components are used to produce multiple products. All components have a common deterministic replenishment lead time (from an uncapacitated supplier) L. 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 goal is to find a control policy, consisting of replenishment and allocation decisions to minimize the total long run average cost, where the instantaneous cost is the sum of linear inventory holding costs and backlog penalty costs.

This is a difficult problem and the optimal policy is not known in general. We approach this problem using a stochastic program (SP) known as a newsvendor network. This is a two stage stochastic linear program with complete recourse. We show that a relaxation of this SP provides a lower bound on the cost achievable in the inventory system under any feasible policy. For two simple examples, the so-called W and M systems, we show how to translate the solution of the SP into a control policy for the inventory system. We introduce and discuss our conjecture that these policies are asymptotically optimal as $L \to \infty$.

[This is joint work with Mustafa Dogru and Qiong Wang.]

Speaker: P. R. Kumar, University of Illinois at Urbana-Champaign
Title: A Formulation and Theory for Delay Guarantees in Wireless Networks
Abstract: Delay guarantees have been problematic in networking. The usual focus of theory is only on providing throughput guarantees. Yet, wireless networks will increasingly need to support applications requiring such guarantees, e.g., voice-over-IP, interactive video, and control over networks. We propose a theoretical framework for addressing the problem of delay guarantees in wireless networks that incorporates three key issues – delay, throughput, and channel reliability – in the specification of quality of service. A somewhat surprising necessary and sufficient condition characterizes when the quality of service requirements of a given set of nodes can be met. It can be checked in nearly linear time, providing a tractable admission control algorithm. Further, there are easily implementable scheduling policies that are feasibility optimal in the sense that they can meet the demands of every feasible set of nodes. The theory can be extended to more general arrival patterns and fading processes, and can also be cast in a utility maximization framework for delay guarantees.

[Joint work with I-Hong Hou and V. Borkar].

References at http://decision.csl.illinois.edu/∼prkumar/html_files/postscript_files.html

Speaker: Kavita Ramanan, Carnegie Mellon University
Title: Reflected Brownian Motions, Dirichlet Processes and Queueing Networks
Abstract: Mike Harrison's pioneering work on heavy traffic limits has spawned over three decades of research on queueing networks and their scaling limits. In particular, it has led to the development of a fairly general theory of semimartingale reflected Brownian motions in non-smooth domains, and their characterization as heavy traffic limits of a broad class of multiclass queueing networks. Motivated by examples of networks arising in applications that lie outside this framework, we provide a potentially useful characterization of a broad class of reflected Brownian motions in quite general domains in terms of so-called Dirichlet processes, and show that this class includes certain non-semimartingale reflected Brownian motions that arise as scaling limits of queueing networks.

[This includes joint work with Weining Kang.]

Speaker: Assaf Zeevi, Columbia University
Title: Stochastic Networks and Parameter Uncertainty
Abstract: The traditional setting in which stochastic processing networks are studied assumes that key problem primitives, e.g., service rates, arrival rates, system load etc are all known to the system manager with perfect precision. Any optimization or performance analysis objective then takes these as an input. A good example is the so-called static planning problem, prominently displayed in many of Mike's papers among others, and plays a central role in the analysis of stochastic networks. The key parameters there are the means of various stochastic primitives, and the solution of this problem dictates/characterizes the "first order" behavior of the network. Based on that, one typically proceeds to study "second order" behavior which is driven by stochastic fluctuations of the primitive processes. Among Mike's contributions to the theory and understanding of stochastic networks lies a lesser known branch that is concerned with parameter uncertainty. The main question being asked there is how does lack of precise knowledge of key problem primitives affect performance analysis and control of the network. For example, suppose that the mean of the arrival process is estimated from past data and hence only its sampling distribution is known rather than a precise value. What is a suitable static planning problem here, and what are the implications on "second order" analysis? While this setting is more representative of what one finds in practice, most of the literature on stochastic networks does not focus on such effects. This talk will provide a brief overview of the subject, and discuss some potential directions to address analysis and control problems in the setting of parameter uncertainty.

[Based in part on joint work with Mike Harrison.]

Speaker: Frank Kelly, University of Cambridge
Title: Brownian Network Models of Ramp Metering
Abstract: Unlimited access to a highway network can, in overloaded conditions, cause a loss of capacity. Ramp metering can help avoid this loss of capacity. The design of ramp metering strategies has several features in common with the design of access control mechanisms in communication networks.

In this talk we describe an approach to the design of ramp metering flow rates informed by results developed for models of Internet congestion control. Our approach is based on Brownian network models that exploit simplifications arising in heavy traffic.

For references see http://www.statslab.cam.ac.uk/∼frank/TALKS/harrison.html