Stochastic Networks Workshop

April 12-13, 2012

Thursday 12 April

9.30		Refreshments and tea
10:15		Opening
10.20	Frank Kelly	Brownian models of congested networks
11.00	Tava Olsen	Heavy traffic control and pricing for systems with leadtime sensitive customers
11.30	Andy Philpott	Supply function equilibrium in electricity transmission networks
12.10		Lunch
2.00	Phil Pollett	The limiting behaviour of a patch occupancy model
2.50	Sophie Hautphenne	Sensitivity analysis of epidemic networks
3.20		Afternoon tea
3.50	Raazesh Sainudiin	Posterior expectation of regularly paved random histograms
4.20	Jeff Hunter	The role of Kemeny's constant in properties of Markov chains
6.30		Conference dinner at Ima, 57 Fort Street

Friday 13 April

9.20	Frank Kelly	Multipath routing and congestion control
10.00	Fernando Beltran	Efficiency comparison of multicast routing cost allocation mechanisms on Internet-like topologies
10.30		Morning tea
11.00	Mark Holmes	Degenerate random environments
12.00		Lunch

All talks take place in Case Room 3, Level 0 of the Owen G Glen Building.

Contact person: Ilze Ziedins ilze@stat.auckland.ac.nz 027 445 2242

This workshop is generously sponsored by the Lion Foundation and the Marsden Fund.

Titles and Abstracts

- **Professor Frank Kelly FRS (University of Cambridge)**
  Brownian models of congested networks

  Brownian models provide tractable high-level descriptions of networks in a variety of application areas. This talk will review work in two areas: the modelling of multi-path routing in the Internet, and the design of ramp metering policies for highway networks.

  In both areas Brownian models are able to exploit the simplifications that arise in heavy traffic, and to make clear the main performance consequences of resource allocation policies.

- **Professor Tava Olsen (University of Auckland)**
  Heavy-traffic control and pricing for systems with leadtime sensitive customers

  This paper studies a queueing model where two customer classes compete for a given resource and each customer is dynamically quoted a menu of price and leadtime pairs upon arrival. Customers select their preferred pairs from the menu and the server is obligated to meet the quoted leadtime. Customers have convex-concave delay costs. The firm does not have information on a given customer's type, so the offered menus must be incentive compatible. A menu quotation policy is given and proven to be asymptotically optimal under traditional large-capacity heavy-traffic scaling.

  Joint work with Baris Ata.

- **Professor Andy Philpott (University of Auckland)**
  Supply function equilibrium in electricity transmission networks

  Many regions in the developed world operate wholesale electricity market pools. In these pools, the system operator takes offers of power from generators at their asking prices and dispatches these offers to meet demand at each location in the transmission grid (represented as a node in a network). The market clearing price at each node in each period is the marginal cost of supply, computed using the asking prices. When the offer curves of generators remain fixed for a number of dispatch periods, or when demand is uncertain, a supply-function equilibrium model is used to predict market outcomes. These models have typically been applied in pools represented by a single node. We present some recent work showing how they might be extended, at least in principle, to markets in transmission networks with many nodes.

  Joint work with Par Holmberg, IFN, Sweden.

- **Professor Phil Pollett (University of Queensland)**
  The limiting behaviour of a patch occupancy model

  We consider a model for the presence/absence of a population in a network of habitat patches, which assumes that colonisation and extinction of patches occur as distinct phases. Since the local extinction probabilities are allowed to vary between patches, our model permits an investigation of the effect of habitat degradation on the persistence of the population. The limiting behaviour of the model is examined as the number n of habitat patches becomes large. We consider two limiting regimes: (i) where the initial number of occupied patches increases at the same rate as n and (ii) where the initial number of occupied patches is almost surely finite in the limit. In case (i) a law of large numbers ensues. In case (ii) our aim is to determine conditions under which a metapopulation that is close to extinction may recover.

  Joint work with Ross McVinish.

- **Dr Sophie Hautphenne (University of Melbourne)**
  Sensitivity analysis of epidemic networks

  In this paper, we provide analytical expressions for the sensitivity of the size of an influenza epidemic spreading among a set of cities connected by air routes in the United States, with respect to two types of parameters: the contamination rates and the travel rates. For that purpose, we approximate the early stages of the epidemic by a multitype branching process, which has the advantage of being more tractable than an epidemic model such as the SIR. We study the resulting changes in the sensitivities when plane contamination is taken into account. Finally, we consider vaccination and analyze the sensitivity of the total size of the epidemic with respect to the fraction of vaccinated people. The sensitivity analysis notably informs us about the impact of small errors in the data on pertinent measures obtained from the model, and highlights the parameters which affect the most the evolution of the disease.

- **Dr Raazesh Sainudiin (University of Canterbury)**
  Posterior expectation of regularly paved random histograms

  We present a novel method for averaging a sequence of histogram states visited by a Metropolis-Hastings Markov chain whose stationary distribution is the posterior distribution over a dense space of tree-based histograms. The computational efficiency of our posterior mean histogram estimate relies on a statistical data-structure that is sufficient for nonparametric density estimation of massive, multi-dimensional metric data. This data-structure is formalized as statistical regular paving (SRP). A regular paving (RP) is a binary tree obtained by selectively bisecting boxes along their first widest side. SRP augments RP by mutably caching the recursively computable sufficient statistics of the data. The base Markov chain used to propose moves for the Metropolis-Hastings chain is a random walk that data-adaptively prunes and grows the SRP histogram tree. We use a prior distribution based on Catalan numbers and detect convergence heuristically by running multiple independent chains. The performance of our posterior mean SRP histogram is empirically assessed for large sample sizes (up to 100 million points) simulated from several high dimensional distributions (up to 1000 dimensions) that belong to the space of SRP histograms.

  If time permits we will mention a novel randomized algorithm that gives the Minimum Distance Estimate under universal L_1 performance guarantees of Devroye and Lugosi from an efficient computation of the Yatracoss class underlying the sequence of SRP histogram states visited by a Markov chain that is driven by a randomised priority queue of its current leaf nodes. Some possible applications to air-traffic management will be illustrated.

  This is joint work with Jennifer Harlow, Dominic Lee and Gloria Teng.

- **Professor Jeffrey Hunter (AUT University)**
  The role of Kemeny's constant in properties of Markov chains

  In a finite m-state irreducible Markov chain with stationary probabilities $\{\pi_i\}$ and mean first passage times m_{ij} (mean recurrence time when $i = j$) it was first shown, by Kemeny and Snell, that $\sum_{j=1}^{m} \pi_j m_{ij}$ is a constant, K, not depending on i. This constant has since become known as Kemenys constant. We consider a variety of techniques for finding expressions for K, derive some bounds for K, and explore various applications and interpretations of these results. Interpretations include the expected number of links that a surfer on the World Wide Web located on a random page needs to follow before reaching a desired location, as well as the expected time to mixing in a Markov chain. Various applications have been considered including some perturbation results, mixing on directed graphs and its relation to the Kirchhoff index of regular graphs.

- **Dr Fernando Beltran (University of Auckland)**
  Efficiency comparison of multicast routing cost allocation mechanisms on Internet-like topologies

  Multicasting is a transmission technique that is increasingly used on the Internet to deliver multicast-based services such as live video. The economics of multicasting is a hard problem to study because of the conflict between Internets best-effort approach and the commercial business objectives of a provider of multicast services. This talk focuses on measuring the efficiency of multicast cost allocation mechanisms on Autonomous System-level network topologies. Efficiency of a cost allocation mechanism for a multicast transmission is defined as the difference between the sum of benefits accrued to users who receive the transmission and the cost of providing the service. A related property is the budget balance condition, that is, the need the network operator has to recover the cost of providing the service. These are conflicting properties: it is not possible to simultaneously achieve efficiency and balance the budget. In this talk, I report on efficiency measures of two multicast routing cost allocation algorithms, the marginal cost and the profit guaranteeing mechanisms, tested on downscaled Internet topologies. The generated topologies closely match real network characteristics, both in terms of node interconnectivity and with respect to two node and link annotations: link latency and membership to Autonomous System. I use the dK-series approach to obtain network topologies by downscaling well-known (real) data sets of annotated Internet topologies. A report is presented on the trade-off between efficiency and budget balance when the algorithms are used to determine cost allocation of multicast routing on the selected networks.

  This is joint work with Daniela Dunn, University of Auckland.

- **Dr Mark Holmes (University of Auckland)**
  Degenerate random environments

  In the middle of each street intersection in a grid-like city, insert a signpost. To each signpost, independently attach a random subset (according to some law) of the 4 possible signs N, S, E, W. The result is a random directed graph (in 2 dimensions) that we call a degenerate random environment. We will discuss the connectivity in such graphs, such as the set of points that you can get to from the origin by following signs, as well as some implications for random walks in random environments.

  Joint work with Tom Salisbury.