Mathematisches Forschungsinstitut Oberwolfach

Report No. 52/2003

Applied Probability

November 30th – December 6th, 2003

This meeting was organized jointly by Francois Baccelli (Paris), Arie Hordijk (Leiden), and Volker Schmidt (Ulm). It continued a sequence of successful Oberwolfach conferences on applied probability, where the previous one was held in 1998.

The present meeting had 46 participants. They came from USA (# 11), Germany (# 10), France (# 5), Netherlands (# 5), Australia (# 2), Denmark (# 2), Great Britain (# 2), Canada, Finland, Greece, Iceland, Israel, Luxembourg, Poland, Russia, Sweden.

The aim of this meeting was to bring together researchers who represent important new developments in the fields of applied probability, in particular

- stochastic networks
- stochastic control problems
- probabilities of large deviations, rare events, and extremes
- spatial stochastic models
- stochastic simulation

There were 9 sessions where actual topics of applied probability were discussed from several viewpoints. Each session started with a longer survey talk followed by a number of shorter communications. These included

- control and optimization in stochastic networks, insurance and finance
- connections between queueing and information theories; their applications to stochastic modelling of communication and computer systems
- multiresolution image analysis and monotonicity issues, in particular for percolation and Gibbs systems such as the Ising and Widom–Rowlinson models
- spatial point processes, germ-grain models, random tessellations, their applications to structural analysis of spatial networks
- importance sampling and differential games
- perfect sampling, rare event simulation with heavy tails, and other recent developments in Monte Carlo simulation
- extremal dependence; applications in data network modelling and finance
- application of fluid models to recurrence and central limit theorems of queueing networks

All participants prepared short reports about their recent research, which were presented on a pin board. This prompted many lively and stimulating discussions. Various new scientific contacts were formed, initiating quite a number of collaborations.

Abstracts

A game theoretic look at a multi-access communication problem

Venkat Anantharam
(joint work with Richard J. La)

We explore the use a cooperative game theoretic formulation of networking problems to understand fair allocation. The program is as follows : (1) determine the characteristic function of the game — this is the function that associates to each subset of the players the maximum they can achieve even if the complementary subset colludes to work against them; (2) define fairness through a reasonable set of axioms – the axioms would be context sensitive; (3) determine which allocations are fair, feasible, and in the core of the game.

The talk describes how we carried out this program for the Gaussian multi-access channel, a basic model for communication over a shared wireless medium. The result of our work is a proposal for how to share the overall Shannon capacity of the channel. The proposed allocation is fair, feasible, and in the core of the game, and it is the unique allocation with these properties. Further, it has many nice intuitive qualitative features as a function of the signal to noise ratios of the individual users of the channel.

Recent developments in rare event simulation with heavy tails

Søren Asmussen
(joint work with D.P. Kroese and R.Y. Rubinstein)

Consider the simulation of $P(Y_1+\cdots+Y_n > u)$ where the Y_i are i.i.d. and heavy-tailed and u is large. Existing i.i.d. importance sampling schemes (Asmussen, Højgaard & Binswanger 2000, Juneja & Shahabuddin 2002) have a considerable freedom in the choice of the twisted distribution. However, this choice is important for the performance. We show here that choosing the twisted distribution as the member of a particular parametric class having minimum entropy distance from the conditional distribution of the Y_i given $Y_1+\cdots+Y_n > u$ gives a precise and efficient suggestion. The implementation details involve maximum likelihood and subexponential asymptotics.

The i.i.d. character of the change of measure in the above scheme is intuitively disturbing because subexponential limit theory predict that one of the Y_i (the largest) has a particular role. We present an asymmetric importance sampling scheme taking this into account as well as a conditional Monte Carlo idea where the estimator is

$$nP(Y_1+\cdots+Y_n > u, M_n = Y_n \,|\, Y_1,\ldots,Y_{n-1}) \;=\; \overline{F}\big(M_{n-1} \vee (u - Y_1 - \cdots - Y_{n-1})\big)$$

where $M_k = \max(Y_1,\ldots,Y_k)$. The schemes are demonstrated numerically to have superior performance compared to existing algorithms.

Portfolio-optimization with jumps and unobservable intensity process

NICOLE BÄUERLE
(joint work with Ulrich Rieder)

We consider a financial market with one bond and one stock. The stock price process follows a geometric Lévy process, where jumps occur according to a Poisson process. Our aim is to maximize the expected CRRA utility of terminal wealth. This problem can be solved explicitly via stochastic control methods and the optimal portfolio strategy invests a constant fraction of the wealth in the stock.

More interesting is the problem when the jumps occur according to a Markov-modulated Poisson process and the investor is not able to observe the driving Markov chain. Using a classical result form filter theory it is possible to reduce this problem with partial information to one with complete information. With the help of a generalized Hamilton Jacobi Bellman equation it is possible to identify an optimal investment strategy.

Approximate decomposition of modulated Poisson-Voronoi tessellations

BARTŁOMIEJ BŁASZCZYSZYN
(joint work with René Schott)

Approximate decomposition is a technique to get approximation formulae for inhomogeneous models by using corresponding results obtained for homogeneous ones. Mathematical formalization of this technique requires to estimate the approximation error.

In this talk we consider the Voronoi tessellation of Euclidian space that is generated by an inhomogeneous Poisson point process whose intensity takes different constant values on sets of some finite partition of the space. Considering the Voronoi cells as marks associated to points of the point process, we prove that the intensity measure (mean measure) of the marked Poisson point process admits an approximate decomposition formula. The true value is approximated by a mixture of the respective intensity measures for homogeneous models, while the explicit upper bound for the remaining term can be computed numerically for a large class of practical examples. By Campbell's formula, an analogous approximate decomposition is deduced for the Palm distributions of individual cells.

This approach makes it possible to analyze a wide class of non-homogeneous Poisson-Voronoi tessellations, by means of formulae and estimates already established for the homogeneous case. Our analysis applies also to the Poisson process modulated by an independent stationary random partition, in which case the error of the approximation of the double-stochastic Poisson-Voronoi tessellation depends on some integrated linear contact distribution functions of the boundaries of the partition elements.

The talk is based on a paper published under the same title in *Advances in Applied Probability* **35** (2003), 847–862.

Subexponentiality and semi-exponentiality. Some new limit theorems for random walks with semi-exponential step distributions

A.A.Borovkov

The talk contains two parts. The first one deals with the characterization of subexponentiality and the relationship between the classes of subexponential and semi-exponential distributions. Zones of additivity of tails are described for semi-exponential distributions. The second part of the talk is devoted to some advanced asymptotic results for random walks with semi-exponential step distribution. The second-order approximation for the distribution of the supremum of sequential sums of random variables and the asymptotics of the probability of large deviations for the first passage time are obtained.

Application of fluid models to recurrence and central limit theorems of queueing networks

Maury Bramson

Over the past decade, fluid models have become one of the main tools for analyzing queueing networks. They can be used to study the stability (i.e., recurrence) and heavy traffic limits (i.e., scaled diffusive limits) of queueing networks. In both cases, this provides a systematic approach for reducing problems in a random setting to simpler deterministic ones. The talk provides a summary of the main results in the area.

Queueing networks with an attached inventory

Hans Daduna

(joint work with Maike Schwarz, Cornelia Sauer, Rafal Kulik and Ryszard Szekli)

In a service system of M/M/1 type for to serve a customer an additional item is needed (e.g. to perform some repair) which has to be taken from an attached inventory. We construct a Markovian model which allows to compute jointly performance measures of the service process and the behaviour of the inventory.

For different inventory management policies (including (r,Q), (r,S) and a random order size policy) and with random nonzero lead times we derive explicit stability conditions. In case of lost sales we obtain explicitly the asymptotic and stationary distribution of the joint (queue length/inventory size) process. It is of product form, i.e. asymptotically the queue length and inventory size fluctuations behave as if they were decoupled.

From this formula we derive the main performance characteristics of the system, e.g., mean queue length, mean inventory, lost sales rate, safety stocks, service levels.

In a further step we incorporate such a server with attached inventory into a classical Jackson or Gordon-Newell network and derive stability conditions and product form steady state as well.

Germ-grain models from the lilypond growth protocol

Daryl J. Daley

The lilypond model, such as a hard-sphere model $\{(x, R(x)) \colon x \in \varphi\}$ in which φ is a locally finite set of points (or, germs) $x \in \mathbf{R}^d$ that are the centers of spherical grains of radius $R(x)$, is realized by the *lilypond growth protocol*: given φ (e.g. a realization of a stationary point process), each germ starts its growth into a grain at the same time and same rate (in all directions); any particular grain stops its growth when for the first time it collides with another grain. Then for all germs x and y, $R(x) + R(y) \le |x - y|$, and for every germ x there is a germ z, a *grain-neighbour* of x, for which the corresponding grains touch so $R(x) + R(z) = |x - z|$.

When φ is a realization of a stationary Poisson process, the volume fraction ϖ_d (i.e. the proportion of space covered by the grains) can be determined analytically when $d = 1$ ($\varpi_1 = 1 - e^{-1}$; see Daley, Mallows and Shepp (2000), *Stoch. Proc. Appl.*) and numerically (e.g. for $d = 2, 3$ in Daley, Stoyan and Stoyan (1999) *Adv. Appl. Probab.*), together with bounds $0 \le 2^d \varpi_d - 1 \le O(d^{1/2})$. The volume V of a randomly chosen grain is known analytically for $d = 1$: $\mathbf{P}\{V > x\} = e^{-2x} \exp(e^{-2x} - 1)$, and its analogues for $d = 2, 3$, standardized to the same mean, are very close. Is there a limit for large d? What if the grains are similarly oriented cubes? What if the initial germ-sets φ come from small i.i.d. perturbations of a lattice like Z^d?

Analysis of the 1-D model with germs located at the points of a renewal process (i.e. including a point at 0) with lifetime d.f. F, proceeds via the integral equation $G(x) = \int_0^x F(x+y)\, G(dy) + [1 - G(x)]F(2x)$ for the d.f. G of the radius of the grain at 0 assuming no grains in $(-\infty, 0-)$.

Importance sampling and differential games

Paul Dupuis
(joint work with Hui Wang)

We consider the problem of design and analysis of importance sampling schemes. The talk has two main points: (1) the asymptotically optimal variance from among a fixed class of importance sampling schemes can often be characterized in terms of a differential game, and (2) the associated Isaacs equation gives a flexible tool for the construction and analysis of schemes. After a brief review of the literature, we outline the connection between importance sampling and games. We first indicate why a game should characterize the optimal variance, and then, in the context of various examples, describe how subsolutions to a related PDE identify useful algorithms. The examples include functionals of Markov chains and an elementary stochastic network.

Maxima on random time intervals of random walks with heavy-tailed increments

Serguei Foss
(joint work with Francois Baccelli, Stan Zachary, and Zbigniew Palmowski)

Consider a random walk $S_0 = 0$, $S_n = X_1 + \ldots + X_n$, $n \geq 1$ with i.i.d. increments having negative mean $-a$ and distribution function F. For $n \geq 1$, let $M_n = \max_{0 \leq k \leq n} S_n$ and let σ be an arbitrary stopping time (which, in particular, may be infinite with positive probability). C. Klüppelberg (1988) introduced the following class $\mathbf{S}^*$ of subexponential distributions:

$$F \in \mathbf{S}^* \quad \text{iff} \int_0^x \overline{F}(u)\overline{F}(x-u)\,du \sim 2m^+\overline{F}(x), \quad x \to \infty.$$

Here $\overline{F}(x) = 1 - F(x)$ and $m^+ = \int_0^\infty \overline{F}(u)\,du$. We proved the following result: $F \in \mathbf{S}^*$ if and only if

$$\mathbf{P}(M_\sigma > x) \sim \mathbf{E}\int_0^\infty \mathbf{P}(\sigma > t)\overline{F}(x + at)dt$$

uniformly for all stopping times σ. Furthermore, we extended the above result to probabilities of crossing non-linear boundaries of the form $x + g(n), n \geq 0$ when $x \to \infty$. Finally, we applied the results to study various queueing networks, including multi-server queues and tandem queues.

Open loop routing in network of queues

Bruno Gaujal

This talk discusses the problem of effective computation of the optimal routing sequence in a queuing system made of two queues.

In the general framework with stationary and ergodic arrivals and services, we show that the optimal open loop routing policy is given by a billiard sequence (generated by the trajectory of a billiard ball on a square table). We also show that the optimal policy (minimizing the expected waiting time) is a Sturmian sequence and we establish several qualitative properties of this policy (monotonicity, continuity, convexity).

For exponential service times and Poisson arrivals, we show how the kernel method can be used to get effective computations of the optimal policy. This is done by using complex analysis and a trick in shifted polar coordinates.

We address the issues of time complexity as well as numerical stability of this algorithm. We then discuss an extensive set of experiments which show several interesting features of the optimal policy with apparent discontinuities and a fractal behaviour.

Recent developments in Monte Carlo simulation

Peter W. Glynn

This talk surveys recent developments in stochastic simulation, with a particular emphasis on the simulation of discrete-event systems. The topics covered include: a large deviations characterization of computational efficiency (analogous to Bahadur efficiency for statistical estimators), a new measure of computational complexity for Monte Carlo schemes, a new class of renewal-theoretic estimators for regenerative processes, Harris (regenerative) structure of discrete-event systems (when viewed as generalized semi-Markov processes (GSMP's)), algorithmic identification of GSMP regenerations, and total variation convergence rates for SDE numerical schemes.

Monotonicity issues in probability theory: a scattered survey

Olle Häggström

Many intriguing problems in probability theory are of the following kind. Consider a stochastic system with a parameter p that governs its dynamics, and some particular aspect A of the dynamics. For instance, p could be a drift parameter of some random motion, and A is the resulting asymptotic speed. Does A depend monotonically on p? Such questions are particularly natural to ask in systems exhibiting critical phenomena or phase transitions.

In this talk, a survey of such monotonicity issues is given. Specific models that are discussed include percolation, Gibbs systems such as the Ising and Widom–Rowlinson models, Monte Carlo Markov chains, and random walks. As to the answers to the main question – whether A depend monotonically on p – we encounter all four of the following answers, the last of which can be quite a surprise:

- Yes, monotonicity holds, and it admits a short and easy proof.
- Yes, it does hold, but the proof is not straightforward.
- The question is open.
- No, the desired monotonicity is demonstrably false.

Limit theorems for functionals of stationary random tessellations

Lothar Heinrich

(joint work with Hendrik Schmidt and Volker Schmidt)

We consider a stationary and ergodic tessellation X_0 given in a convex, compact sampling window $W_\varrho \subset R^d$. The cells of X_0 possess random inner structures (e.g. point patterns, fibre systems, germ-grain models, tessellations etc.) generated independently of each other and independent of the initial tessellation X_0 by a generic random set model which is associated with a vector $J_0(.)$ of stationary random measures each component acting on R^d. We study the asymptotic behaviour (as $W_\varrho \uparrow R^d$) of some multivariate random functional which is determined by X_0 and the individual cell structures contained in the sampling window W_ϱ. It turns out that this functional provides an unbiased estimator of the intensity vector defined by the generic random vector measure $J_0(.)$. Furthermore, the asymptotic covariance matrix of the functional is calculated and shown to depend only on the distribution of the typical cell of X_0 and the second-order moment and cross covariance measures of the vector measure $J_0(.)$.

Under natural restrictions using the ergodicity of X_0 and the conditional independence between distinct cell structures we prove a strong law of large numbers and a multivariate central limit theorem of the normalized functional.

Finally, some numerical examples are discussed in detail, for which the inner structure of the cells of X_0 is induced by iterated Poisson-type tessellations. The obtained results allow to construct flexible models of random tessellations and make them tractable for a statistical analysis including model fitting. An application to stochastic subscriber line models is briefly discussed.

Perfect sampling: techniques and challenges

Mark Huber

Perfect sampling algorithms generate random variates from probability distributions with an unknown normalizing constant. Protocols such as Coupling From the Past (CFTP) and the Randomness Recycler (RR) employ very different approaches to building such algorithms. CFTP emphasizes complete coupling of Markov chains. RR creates a family of distributions forming a chain between a distribution that is easy to sample from and the target distribution. Ideas such as monotonicity and bounding chains can be used with CFTP, but not RR. Open questions include how to adjust these methods to work with both protocols, and developing new protocols to take advantage of properties of Markov chains related to rapid mixing (such as the spectral gap or conductance methods).

Multiresolution image analysis and Ising models

Wilfrid Kendall

(joint work with Roland Wilson)

This talk describes results of an investigation into percolation and Ising models defined on generalizations of quad-trees used in multiresolution image analysis. These can be viewed as trees for which each mother vertex has 2^d daughter vertices, and for which daughter vertices are linked together in d-dimensional Euclidean configurations. Retention probabilities / interaction strengths differ according to whether the relevant bond is between mother and daughter, or between neighbors. Bounds are established which locate phase transitions and show the existence of a coexistence phase for the percolation model. Quantitative results are extended to the corresponding Ising model using the Fortuin-Kasteleyn random-cluster representation, and compared with computer simulations.

The Skorokhod reflection problem on partially ordered sets

Takis Konstantopoulos

(joint work with Venkat Anantharam)

One formulation of the Skorokhod reflection problem with two boundaries can be stated in terms of ordering. So it is natural to consider the domain of the functions to be a partially ordered set with a least element. In doing so, we get rid of continuity assumptions, usually required for proving existence and uniqueness of the solution. By using order arguments, we show that, in a very general setup, the solution exists, is unique, and satisfies a fixed point equation. A particular case of interest is when the two boundaries coincide. In this case, the problem becomes that of the minimal decomposition of a function, defined on a partially ordered set, as the difference of two increasing functions.

A case of interest is when the poset is a σ-algebra of subsets of a set. Then, the Skorokhod reflection of a signed measure, when the two boundaries are set to zero, is proved to be equivalent to the Hahn-Jordan decomposition of the signed measure; when the boundaries are non-identical (positive) measures we obtain a generalization of the Hahn-Jordan decomposition: indeed, the reflectors turn out to be mutually singular σ-additive functions, i.e., they inherit the property of the signed measure.

Other cases of interest are: (i) The poset of probability distributions on a finite set ordered in the majorization order. Here, increasing (decreasing) functions are the Schur convex (concave) functions. We discuss the reflection problem when the two boundaries are zero. (ii) The poset of positive semi-definite matrices ordered as: $A \leq B \Leftrightarrow B - A$ is positive semi-definite. Here, natural increasing functions are the eigenvalues of a matrix. We thus seek the minimal decomposition (solution to the Skorokhod reflection problem with boundaries identically zero) of certain functions of the eigenvalues. We also reflect on the fact that generalizing the problem from a linearly ordered set to a partially ordered one not only generates several interesting cases, but also sheds new light to the classical problem.

Monotonicity and comparison of tandem systems

Ger Koole

Let $V : N^m \to R$ be a solution of the dynamic programming optimality equation of a certain m-dimensional queueing system. Monotonicity properties of queueing systems can be seen as properties of the corresponding optimality equation or value functions (which converge to the optimality equations).

There are basically two types of convexity results for multi-dimensional queueing systems. The stronger one, multimodularity, holds for a very limited class of queueing systems. Essentially the only system in this class is a tandem queueing system consisting of single-server queues or servers with a variable rate that are controlled optimally (Weber & Stidham 87). Only for two dimensions this class is much more rich, containing for example also routing problems (Hajek 84).

An important class of problems for which multimodularity does not hold are multi-server queues. For this class of problems we obtained results that are weaker than multimodularity: directional convexity. Directional convexity of a function $f : N^m \to R$ means that

$$f(x + y) + f(x + z) \leq f(x) + f(x + y + z) \tag{1}$$

for all x, y and $z \in N^m$. It is equivalent to Equation (1) with y and z restricted to unit vectors. If, in this case, $y = z$ then we get componentwise convexity; if $y \neq z$ then (1) becomes supermodularity. Thus directional convexity is equivalent to componentwise convexity and supermodularity.

Directional convexity has been very useful in obtaining comparison results. It has been used for example in Altman et al. 01 to show that fluid limits of tandem queueing systems have a lower workload than the original queueing systems. Unfortunately directional convexity is of less help when it comes to proving properties of optimal policies. In conclusion, generalizing the Hajek-model to multi-server queues remains an open question.

The lilypond growth protocol, descending chains and invariance properties of Palm distributions

Günter Last

(joint work with Daryl Daley)

We consider a stationary point process N in $\mathbb{R}^d$ satisfying a non-lattice type condition. A descending chain (in N) is an infinite sequence of mutually different N-points such that the distance between two consecutive points is decreasing. We find analytic conditions implying that N is almost surely descendant chain free. Examples are Cox, Poisson cluster and Gibbs processes satisfying certain (exponential) moment conditions. Descendant chain free point processes are then shown to have interesting properties as the absence of percolation in the associated lilypond model (a specific hard-sphere model built on N) and the existence of bijective point maps the leave the Palm distribution of N invariant.

A queue and a store

Jean Mairesse
(joint work with Moez Draief and Neil O'Connell)

The main purpose is to clarify the interplay between two models of queueing theory. The first model is the M/M/1 queue. The second model, less common but very natural, can be described as a queue operating in slotted time with batch arrivals and services. To clearly distinguish between the two models, we describe the second one with a different terminology and as a *storage model*.

The M/M/1/∞/FIFO queue (or M/M/1 queue) is the queue with a Poissonian arrival stream, exponential services, a single server, an unlimited buffer capacity, and a First-In-First-Out service discipline. It can be argued that the M/M/1 queue is the most elementary and the most studied system in queueing theory. Let $\mathbf{A}$ be the process of arrivals and let $(s_n)_{n\in\mathbf{Z}}$ be the sequence of service times of the customers. We call them *input variables* and we call *output variables* the process of depatures from the queue $\mathbf{D}$ and the sequence $(r_n)_{n\in\mathbf{Z}}$, where r_n is the time spent by customer n at the *very back* of the queue. Consider now the storage model. The equations driving the dynamic are exactly the same as for the single server queue. But it is not the same variables that make sense in the two models. The important variables are the ones corresponding to the departures from the system. The departures are coded in the variables $(D_n)_n$ for the single server queue and in the variables $(r_n)_n$ for the storage model.

This duality between the two models gets reflected at different levels. At the level of the equilibrium behaviour: we prove that the output variables $(\mathbf{D}, r)$ have the same distribution as the input variables $(\mathbf{A}, s)$. This holds when the process $\mathbf{A}$ is Poisson and the r.v.'s $(s_n)_n$ are exponentially distributed, or when the process $\mathbf{A}$ is Bernoulli and the r.v.'s $(s_n)_n$ are geometrically distributed (Burke type theorems). In the geometric case, we propose a new way of proving this theorem by using the *zigzag process*, a symetrized version of the workload process.

A second facet of the duality between queues and stores appears when studying the transient behaviour of a tandem by means of the Robinson-Schensted-Knuth algorithm: the first and last row of the resulting semi-standard Young tableau are respectively the last instant of departure in the last queue and the total number of departures from the last store. [This builds on work by O'Connell. See for instance: N. O'Connell. A path-transformation for random walks and the Robinson-Schensted correspondence. *Trans. Amer. Math. Soc.*, 355:3669–3697, 2003.]

Comparing strength of locality of reference: popularity, majorization, and some folk theorems for the miss rate and output of caches

Armand M. Makowski
(joint work with Sarut Vanichpun)

The performance of demand-driven caching is known to depend on the locality of reference exhibited by the stream of requests made to the cache. Yet, like the notion of burstiness used in traffic modelling, locality of reference, while endowed with a clear intuitive content, admits no simple definition. In spite of numerous efforts, no consensus has been reached on how to formalize this notion, let alone on how to compare streams of requests on the basis of their locality of reference.

We take on this issue with an eye towards validating operational expectations associated with the notion of locality of reference. More specifically, we focus on two "folk theorems," namely (i) The stronger locality of reference, the smaller the miss rate of the cache; (ii) Good caching is expected to produce an output stream of requests exhibiting less locality of reference than the input stream of requests.

We discuss results concerning these two folk theorems in the context of a cache operating under a demand-driven replacement policy when document requests are modelled according to the Independent Reference Model (IRM). As we propose to measure strength of locality of reference in a stream of requests through the skewness of its popularity distribution, we introduce the notion of majorization as a means of capturing this degree of skewness.

We show that these folk theorems hold for caches operating under a large class of cache replacement policies, including the optimal policy A_0 and the random policy, but may fail under the LRU and CLIMB policies. In addition, we explore how the majorization of popularity distributions translates into comparisons of three well-known locality of reference metrics, namely the inter-reference time, the working set size and the stack distance.

Broadcasting on trees and the hard-body model

James Martin

We consider a branching random walk with binary state space and index set $\mathbb{T}^k$, the infinite rooted tree in which each node has k children (also known as the model of *broadcasting on a tree*). The root of the tree takes a random value 0 or 1, and then each node passes a value independently to each of its children according to a 2×2 transition matrix $\mathbf{P}$. We say that *reconstruction is possible* if the values at the dth level of the tree contain non-vanishing information about the value at the root as $d \to \infty$. Adapting a method of Brightwell and Winkler, we obtain new conditions under which reconstruction is impossible, both in the general case and in the special case $p_{11} = 0$. The latter case is closely related to the *hard-core model* from statistical physics; a corollary of our results is that, for the hard-core model on the $(k+1)$-regular tree with activity $\lambda = 1$, the unique simple invariant Gibbs measure is extremal in the set of Gibbs measures, for any k.

Exact asymptotics of queueing networks

David R. McDonald
(joint work with Robert D. Foley)

We extend the Markov additive methodology developed in [2, 4] to obtain the sharp asymptotics of the steady state probability of a queueing network when one of the nodes gets large. We focus on a new phenomenon we call a bridge. The bridge cases occur when the Markovian part of the *twisted* Markov additive process is one null recurrent or one transient while the jitter cases treated in [2, 4] occur when the Markov additive process is (one) positive recurrent. The asymptotics of the steady state is an exponential times a polynomial term in the bridge case but is purely exponential in the jitter case.

This bridge phenomenon is ubiquitous. We review the bathroom problem described in [6] and find the bridge case is present. We derive the sharp asymptotics of the steady state distribution of the number of customers queued at each node as the number of customers queued at the Server One grows large. In doing so we get an intuitive understanding of the flat boundary theory in [6]. We obtain the same results obtained in [1] using complex variable methods.

Finally let's say a word about the matrix-geometric method. The level is the additive component and the phase is the Markovian component of a Markov additive process with kernel K^∞. For a QBD process the additive component is nearest neighbor. Denote the transition probabilities from phase i at level ℓ to phase j at level $\ell-1$, respectively ℓ and $\ell+1$ by A, respectively B and C. Hence we can represent the Feynmann-Kac kernel as

$$\hat{K}_\gamma(i,j) = \sum_t e^{\gamma t} K((i,0);(j,t)) = (\nu^{-1}A + B + \nu C)_{ij}.$$

if $\nu = \exp(\gamma)$. Solving the Riccatti equation associated with the matrix-geometric method is equivalent to finding a Perron-Frobenius eigenvector of the Feynmann-Kac kernel with eigenvalue one. That means solving $(\nu^{-1}A + B + \nu C)\mathbf{y} = \mathbf{y}$ is equivalent to finding γ and $\hat{h} = \mathbf{y}$ so $\hat{K}_\gamma \hat{h} = \hat{h}$. Equivalently, the solution to the Riccatti equation gives a harmonic function $h(x,i) = \exp(\gamma x)\hat{h}(i)$ for K^∞.

References

[1] Flatto, L., Hahn, S. (19784). Two Parallel Queues Created by Arrivals with Two Demands I, *SIAM J. Appl. Math,* **44**, 1041-1053.

[2] Foley, R. and McDonald, D. (2001). Join the Shortest Queue: Stability and Exact Asymptotics, *Ann. Appl. Probab.* **11**, 569–607.

[3] Foley, R. and McDonald, D. (2004). Large deviations of a Queueing Network with a Cross-Trained Server, *submitted Ann. Appl. Probab.*

[4] McDonald, D. (1999). Asymptotics of first passage times for random walk in a quadrant, *Ann. Appl. Probab.* **9**, 110–145.

[5] Neuts, M. F. (1981). *Matrix-Geometric Solutions in Stochastic Models; an Algorithmic Approach,* The Johns Hopkins University Press, Baltimore.

[6] Shwartz, A. and Weiss, A. (1994). *Large deviations for performance analysis,* Chapman and Hall.

Value functions, performance evaluation, and optimization in stochastic networks

Sean Meyn

This paper concerns control and performance evaluation for stochastic network models. Structural properties of value functions are developed for controlled Brownian motion (CBM) and deterministic (fluid) workload-models. Based on these results, the following conclusions are obtained: Outside of a null-set of network parameters,

(i) The fluid value-function is continuously differentiable. Under further minor conditions, the fluid value-function satisfies the derivative boundary conditions that are required to ensure it is in the domain of the extended generator for the CBM model.

(ii) The fluid value-function provides a shadow function for use in simulation variance reduction for stochastic models. The resulting simulator satisfies an exact large deviation principle, while a standard simulation algorithm does not satisfy any such bound.

(iii) The fluid value-function provides upper and lower bounds on performance for the CBM model. This follows from an extension of recent linear programming approaches to performance evaluation.

Internet graphs and Internet queues

Ilkka Norros

(joint work with Hannu Reittu, Petteri Mannersalo, Michel Mandjes, Miranda van Uitert)

Two different topics are discussed in this talk. First, the so called Power-Law Random Graph (PLRG) is considered, where the degrees of the nodes are drawn independently from a distribution with power tail, and the connections are then chosen randomly. Denote the number of nodes by N and assume that the degree distribution is $\mathrm{P}(D > d) = d^{-\tau+1}$ with $\tau \in (2, 3)$ (infinite variance!). Then, the PLRG has a giant component, and the hop distance between two randomly selected nodes of the giant component is, asymptotically almost surely, less than $(2 + o(1))k^*$, where

$$k^* = \left\lceil \frac{\log \log N}{-\log(\tau - 2)} \right\rceil .$$

Thus, the practical diameter of the graph grows hardly at all when N runs through many orders of magnitude. This feature is based on a kind of "soft hierarchy": the nodes with high degrees form a "core network" that enables short connections. The core can be further split into $k^* + 1$ disjoint "layers" such that each node in a layer is a.a.s. directly connected to some node in a strictly higher layer.

The second topic is a problem in queueing systems with Gaussian input. Let $Z = (Z_t)_{t\in\Re}$ be a continuous centered Gaussian process with stationary increments, and let R be its reproducing kernel Hilbert space. Consider the event $B = \{Z_t \geq t \ \forall t \in [0, 1]\}$. Assume that Z has non-differentiable paths. Then, there is a closed set $S^* \subseteq [0, 1]$ such that the most probable path in B (i.e., the path with minimal R-norm) is

$$\beta^*(t) = \mathrm{E}\left[Z_t \,|\, Z_s = s \ \forall s \in S^* \right].$$

In the case that Z is a fractional Brownian motion with self-similarity parameter $H \in (0, 1)$, the set S^* can be characterized as follows:

$$S^* = \begin{cases} [s^*_H, 1], & \text{if} \quad H < 1/2 \\ \{1\}\,, & \text{if} \quad H = 1/2 \\ [0, s^*_H] \cup \{1\}\,, & \text{if} \quad H > 1/2, \end{cases}$$

where $s^*_H \in (0, 1)$. We also have analytical expressions for β^* and $\|\beta^*\|$ in terms of s^*_H.

Identifying Markov chains with a given invariant measure

Phil Pollett

Let $Q = (q_{ij},\ i, j \in S)$ be a stable and conservative q-matrix of transition rates over a countable set S. Suppose that we are given a subinvariant for Q, that is, a collection of positive numbers $m = (m_i,\ i \in S)$ such that $\sum_{i\in S} m_i q_{ij} \leq 0,\ j \in S$. Our problem is to identify a Q-process for which m is invariant, that is, a standard transition function $P(\cdot) = (p_{ij}(\cdot),\ i, j \in S)$ that satisfies $p'_{ij}(0+) = q_{ij}$, $i, j \in S$, and

$$\sum_{i\in S} m_i p_{ij}(t) = m_j, \qquad j \in S,\ t > 0.$$

We begin by showing that if m is invariant for P, then it is subinvariant for Q, and then *invariant* for Q if and only if P satisfies the backward differential equations. A simple corollary is that if m is invariant for *minimal* Q-process, then it is invariant for Q.

The major result gives conditions for the existence of a Q-process P for which the given measure m (subinvariant for Q) is invariant for P; one such Q-process is specified through its resolvent. The invariance condition is shown to be necessary in the case where Q is single-exit (there is a single escape route to infinity). We also give necessary and sufficient conditions for this process to be *honest*, that is $\sum_{j\in S} p_{ij}(t) = 1$ for all $i \in S$ and $t > 0$, as well as a simple sufficient condition for the existence of an honest Q-process for which the given measure m is invariant for P. The case where Q is symmetrically reversible with respect to m is considered in some detail, leading to a complete solution of the existence and uniqueness problem for birth-death processes.

Connections between queueing and information theories

Balaji Prabhakar

(joint work with Robert Gallager, Chandra Nair, and Devavrat Shah)

A collection of theorems in point process theory assert that the Poisson process results as the limit of repeatedly performing certain operations on an arbitrary initial (input) process. These operations include, for example, the splitting of a point process, the superposition of independent copies of a point processes, random translations, and queueing through an exponential server queue. Since the Poisson process has the highest entropy rate of all processes at a given rate, one is naturally lead to the question: Do these operations increase the entropy of the input process?

In this talk we show that some queueing systems indeed increase the entropy. There are two main ideas in our argument: one queueing-theoretic, and the other information-theoretic. The queueing-theoretic idea consists of viewing the evolution of queueing processes in the reverse direction of time and comparing the entropy of the variables involved the forwards and reverse directions. The information-theoretic idea is to consider the "mutual information" between the input (arrival process) and the output (departure process) as a way of obtaining a key equation relating the entropy of the input to that of the output.

Some useful by-products are discrete-time versions of: (i) a proof of the celebrated Burke's Theorem, (ii) a proof of the uniqueness, amongst renewal inputs, of the Poisson process as a fixed point for exponential server queues, (iii) connections with the timing capacity of queues.

Extremal dependence: applications in data network modelling and finance

Sidney Resnick

For multivariate heavy tailed phenomena, extremal dependence analysis assesses the tendency of large values of components of a random vector to occur simultaneously. This kind of dependence information can be qualitatively different than what is given by correlation which averages over the total body of the joint distribution. Also, correlation may be completely inappropriate for heavy tailed data.

We review some techniques, somewhat exploratory in nature, for assessing asymptotic independence. Examples are given: (a) The vector of (internet file size, throughput, duration of transfer); (b) The vector of exchange rate returns relative to the dollar (prior to introduction of the Euro) of (French Franc, German Mark, Japanese Yen). In an attempt to formalize a procedure, we introduce a summary measurement called the extremal dependence measure (EDM), a measure of the tendency of large values of components of a random vector to occur simultaneously and show consistency and asymptotic normality properties for the standard case of multivariate regular variation. Initial experiments with the EDM are promising. We also discuss a subfamily of distributions possessing asymptotic independence call the hidden regularly varying class.

Asymptotic analysis of controllable queueing networks

Ad Ridder
(joint work with Arie Hordijk and Alexander Gajrat)

We consider a stochastic tandem queueing network in discrete-time with control of the service rates in order to minimize expected total holding cost upto a finite horizon. This model can be formulated as a Markov decision problem, which is not trivally solved for certain parameter settings, namely when the service rate at the first queue is larger than that at the second queue, while the holding cost at the first queue is less than at the second queue. The associated deterministic fluid control model has a simple fluid solution, so the idea is to find policies (in the MDP) for which the fluid limits of the stochastic trajectories coincide exactly with the optimal fluid paths. These policies are called asymptotically fluid optimal (a.f.o.). It has been shown that in this particular model the a.f.o. policies are determined by a family of sub-linear switching curves.

Our next issue is finding "the best" switching curve within this family. When analyzing the asymptotics of the value function associated to an a.f.o. switching-curve-policy in the stochastic problem, we derive an expression consisting of more terms. The first term gives always the value of the optimal fluid solution; the second term may be used to discriminate between a.f.o. policies. It results in a specific logarithmic switching curve that would be "the best" a.f.o. policy. For the analysis of the asymptotics of the value function we apply large deviations results for random walks.

Markov processes conditioned to never exit a subspace of the state space with application to the single server queue

Tomasz Rolski
(joint work with Zbigniew Palmowski)

The theory of Markov processes, never exiting a subspace of the state space is presented. Different equivalent approaches are considered: as a limiting conditional process or by the change of the probability measure via exponential martingales. We show the relationship between exact asymptotics of the exit time from a subspace and the exponential martingale used to change the measure. The special case of piecewise deterministic Markov processes (PDMP) is worked out. Examples from the theory of single server queues are presented, in particular, within the framework of PDMP's, we characterize the workload process in M/G/1 queue conditioned to stay positive.

Control of ruin probabilities by discrete-time investments

Manfred Schäl

The Cramer-Lundberg insurance model is studied where, in addition, the risk process can be controlled by investments in a risky asset. The price process of the asset is driven by a compound Poisson process. The performance criterion is the ruin probability. The problem will be imbedded in the framework of discrete-time stochastic dynamic programming (Markov decision processes) for minimizing costs. Basic tools for finding good investment plans are the verification theorem and the Howard improvement (where the latter concept is not defined for continuous–time control).

Variograms and their analogies for spatial extremes and in marked point process theory

Martin Schlather

(joint work with Peter Diggle, Tilmann Gneiting, Paulo Ribeiro, Zoltán Sasvári and Jonathan Tawn)

The aim of this research work has been twofold. First, the properties of the variogram and the covariance function for random fields has been investigated in more detail. Second, analogue notions to the variogram or the covariance function have been defined for other spatial processes, and the properties of these characteristics have been examined. Our focus has been on spatial extreme value processes and marked point processes.

Let $Z(x)$, $x \in \mathbf{R}^d$, be a second-order stationary random field, $C(h) = \mathrm{cov}(Z(x), Z(x+h))$, $h \in \mathbf{R}^d$, the covariance function and $\gamma(h) = \frac{1}{2}\mathbf{E}(Z(x) - Z(x+h))^2$, $h \in \mathbf{R}^d$, the variogram. Several relations hold between the set of covariance functions and the set of variograms in $\mathbf{R}^d$. For example, if γ is a bounded variogram and $\gamma(h) = C(0) = C(h)$ holds, then C is a covariance function if and only if

$$C(0) \geq \lim_{u\to\infty} (2u)^{-d} \int_{[-u,u]^d} \gamma(h)\,\mathrm{d}h.$$

If Z is a stationary max-stable random field with unit Fréchet margins, then the variance of Z does not exist. Here, the extremal coefficients function θ,

$$\theta(h) = \lim_{x\to\infty} \frac{\log(\mathbf{P}(Z(h) \leq x, Z(0) \leq x))}{\log \mathbf{P}(Z(0) \leq x)}, \qquad h \in \mathbf{R}^d,$$

can be defined where $\theta(\cdot) - 1$ is again a variogram. The function is not differentiable at the origin except θ is identically constant. Hence, the ensemble of extremal coefficients functions on $\mathbf{R}^d$ can be identified with a subclass of the set of variograms on $\mathbf{R}^d$.

In case of a marked point process Ψ, the geostatistical characteristics can be redefined as conditional quantities given there are points of the corresponding unmarked point process at the locations under consideration. The precise definition can be based on Palm distribution or, directly on Radon-Nikodym derivatives for suitable σ-algebras. However, here, the mark variogram is neither necessarily conditionally negative definite, nor the relation $\gamma(h) = C(0) = C(h)$ must hold.

Optimization problems in non-life insurance

Hanspeter Schmidli

We consider a classical Cramér–Lundberg risk model where the aggregate claim follows a compound Poisson process. In addition we allow (proportional) reinsurance and investment into a risky asset modelled as a Black–Scholes model. Then the surplus process fulfils

$$\mathrm{d}X_t^{A,b} = (c(b_t) + mA_t)\,\mathrm{d}t + \sigma A_t\,\mathrm{d}W_t - b_{t-}\,\mathrm{d}S_t\ ,$$

where S_t is the aggregate claims process. The goal is to maximize the survival probability $\mathbb{P}[\inf_{t\geq 0} X_t \geq 0]$. The Hamilton–Jacobi–Bellman equation connected to the problem is

$$\sup_{A,b} \tfrac{1}{2}\sigma^2 A^2 f''(x) + (c(b) + mA)f'(x) + \lambda\Big[\int_0^{x/b} f(x - by)\,\mathrm{d}G(y) - f(x)\Big] = 0\ .$$

We first show that a twice continuously differentiable solution to the HJB equation exits. Then we prove a verification theorem, showing that any solution to the HJB equation is a multiple of the value function, and that an optimal strategy can be found by choosing (A_t, b_t) as the argument maximizing the HJB equation at $x = X_t$.

In the small claim case (exponential moments exist) we prove that the optimal ruin probability $\psi^*(x)$ behaves like $\zeta\, \mathrm{e}^{-Rx}$, where R is the largest adjustment coefficient for constant strategies (A, b). Let (A^*, b^*) denote the arguments that maximize the adjustment coefficient. Then the optimal strategy (A_t, b_t) converges to (A^*, b^*).

In the large claim case (no exponential moments exist) we can also define the largest adjustment coefficient R. Here $b^* = 0$. It turns out that $\psi^*(x)\, \mathrm{e}^{Rx}$ is a non-increasing function, and therefore converges. If the distribution tail of the claims is heavy enough we find that the limit is positive. Also in this case the optimal strategy (A_t, b_t) converges to $(A^*, 0)$.

Response times in M/M/s fork-join networks

Richard Serfozo
(joint work with Sung-Seok Ko)

We study a fork-join processing network in which jobs arrive according to a Poisson process, and each job splits into m tasks, which are simultaneously assigned to m nodes that operate like $M/M/s$ queueing systems. When all of its tasks are finished, the job is completed. We present a closed-form formula for approximating the distribution of the network's response time (the time to complete a job) in equilibrium. We also present an analogous approximation for the distribution of the equilibrium queue length (the quantity of jobs in the system), when each node has one server. Statistical tests show that these formulas are good fits for the distributions obtained from simulations.

Control of queueing networks in the large deviations domain

Adam Shwartz
(joint work with Rami Atar, Paul Dupuis, and Nahum Shimkin)

Asymptotic analysis of control problems for large and fast stochastic networks leads to control problems in the large deviation scaling. This is not well defined since under most control policies it is not known whether large deviations estimates hold. We propose two formulations that lead to meaningful models. These are but the first steps in developing a comprehensive theory.

1) *Asymptotic escape time criterion.* We consider a network with J queues, with possible arrivals to each, and K servers, each responsible for some of the queues (without overlap). The control is to choose which queue is served by each server at any given time. The usual large deviations scaling is used, namely X^n is the vector of queue sizes for the process whose rates were accelerated by n, while each arrival/departure changes the state vector by $\pm 1/n$. Let σ^n denote the exit time of X^n from a given set G in $\mathbf{R}^J$. For $c > 0$ and $x \in G$ define

$$V^n(x) := -\inf n^{-1} \log \mathbf{E}_x^{u,n} e^{-c\sigma^n},$$

where the infimum is over all controls u. Our goal is to maximize the exit time σ^n. This escape time criterion is useful e.g. in regulating the occurrence of buffer overflows. Our risk-sensitive escape time criterion heavily penalizes early exits. The properties of this problem are studied in the large buffer limit, and related to the value of a deterministic differential game with constrained dynamics. Roughly speaking, this game is a continuous time, deterministic game where one player chooses the control (that is, which queue to serve), while the other player tries to minimize the escape time by skewing the rates. The dynamics are then described by a differential equation, while the value has the following form. If player 1 choose the function $u(t)$ as its control and player 2 choose $m(t)$ then the cost starting at x is

$$C(x,u,m) := \int_0^\sigma [c + \rho(u(t),m(t))]\ dt.$$

For the form of ρ and the details of this formulation see [1]. The difficulty in analyzing this game is that the dynamics are constrained so that all coordinates are kept non negative. We prove that the game has a value V, which is the (viscosity) solution of a PDE. Moreover, we showed that under mild conditions on G (which hold for rectangles),

$$\lim_{n\to\infty} V^n(x) = V(x) \qquad \text{uniformly in } x \in G.$$

For two simple networks, the value is computed explicitly [1, 2], demonstrating the applicability of the approach.

2) *Decay rate for controlled MDP's.* The setup here is of a discrete time Markov decision process. Let x_n, u_n denote the control and action at time n, and let c be the immediate cost function. Define the sample average cost

$$\overline{X}_n := \frac{1}{n}\sum_{i=0}^{n} c(x_n, u_n).$$

We consider the average cost functional, and define an undesirable set B of values and attempt to identify policies which are the least likely (in the large deviations sense) to achieve this set of values. Namely define

$$I_{min}(B) := -\limsup_{n\to\infty} \frac{1}{n} \log \max_\pi \mathbf{P}^\pi \left(\overline{X}_n \in B\right)$$
$$I_{max}(B) := -\liminf_{n\to\infty} \frac{1}{n} \log \min_\pi \mathbf{P}^\pi \left(\overline{X}_n \in B\right).$$

Preliminary results show that there is a gap between these two values, that under some conditions they are achieved by stationary policies, but I_{min} is achieved only by a randomized policy. The conditions on B are required for the lower bound only (I_{max}), where B should be convex. For non convex sets, counterexamples show that optimal policies are history dependent.

References

[1] Atar, R., P. Dupuis and A. Shwartz, "An excape time criterion for queueing networks: asymptotic risk-sensitive control via differential games," *Mathematics of Operations Research*, to appear. www.ee.technion.ac.il/~adam/PAPERS/RiskSens-DiffGames.pdf.

[2] Atar, R., P. Dupuis and A. Shwartz, "Explicit solutions for a network control problem in the large deviations regime," *Queueng Systems: Theory and Application*, to appear. CCIT Report 401, Technion, 2002. www.ee.technion.ac.il/~adam/PAPERS/tandem.ps

[3] Shimkin, N. "Extremal large deviations in controlled i.i.d. processes with applications to hypothesis testing," *Advances in Applied Probability* **25** (1993), 875–894.

Scattering and cycles for face-homogeneous random walks

Flora Spieksma
(joint work with Arie Hordijk and Nicolai Popov)

So-called face-homogeneous random walks are discussed. These are time-homogeneous Markov chains in discrete time, moving on a countable state space $\mathbf{S} = \mathbb{Z}^{p-l} \times \mathbb{Z}_{+}^{l}$. The state space has a finite partition, within each subset of which the walk behaves as a homogeneous random walk.

We are interested in the long run behaviour of this walk: how to determine whether it is positive recurrent, and if transient, which are the escape directions to infinity and how do the scattering probabilities over the different escape directions depend on the initial state?

The motivation for studying this type of walk comes from queueing theory. The study of escape directions may be important for identifying bottleneck nodes or queues.

Denote by $\xi_n(x)$ the position of the walk at time n given initial position x. One method for studying the posed questions, is by studying the fluid limit $\xi_{[tN]}(x_n)/N$ (provided it exists), where x_N may be a fixed point x or a scaled initial point $[xN]$.

By face-homogeneity, stochasticity may be reduced by considering the scaled accumulated mean drift process instead. Even though the limit may not exist, as $N \to \infty$, for many models one can derive properties of the set of possible limits $u(x;t)$, $t \geq 0$, for $x = \lim_{N\to\infty} x_N/N$. These are called fluid paths. They are piecewise linear, continuous paths with right derivative $d^{+}u(x;t)/dt$ called the fluid drift at $u(x;t)$.

Lyapunov function techniques based on the fluid paths are known to work for determining first positive recurrence or transience of the walk. But also, in case of transience, they can be used to determine the almost closed set structure of the walk, introduced by Blackwell. It is intuitively obvious that if no fluid path contains a cycle, then the state space consists of finitely many almost closed atoms. If the latter holds, then we conjecture that to each atom corresponds precisely one fluid path escaping to infinity. The conjecture has been checked for some low-dimensional examples.

The collection of cyclic fluid paths are conjectured to correspond to a completely non-atomic almost closed set. In this case, considering the fluid limit has no meaning at all, since this limit cannot exist for a fixed initial state. Intuitively, this is due to the fact that almost closed sets are not invariant under scaling. The right object to consider, should be the scattering probabilities over the cyclic paths, which stand in one-one correspondence to the almost closed sets. This conjecture has been partially proved for a simplest version of a face-homogeneous random walk with cyclic fluid paths.

A new method for structural analysis of spatial networks

Evgueni Spodarev
(joint work with Simone Klenk and Volker Schmidt)

In applications, one often tackles the problem to distinguish automatically between two images with similar structure. Examples range from testing the structure of road networks, via coverage properties of mobile communication networks to automatic cancer diagnostics in medicine. In this talk, a new method for such analysis is discussed using tools of convex and stochastic geometry. The core of the method is the statistical estimation of morphological image characteristics.

First, a binary image in $\mathbb{R}^d$ is interpreted as a realization of a spatially homogeneous random closed set belonging to the extended convex ring. Second, the vector of all specific Minkowski functionals comprising the volume fraction, the specific surface area and the Euler–Poincaré characteristic (porosity) is computed and its asymptotic properties are studied. This vector is a solution of a system of linear equations resulting from the generalized Steiner formula on the convex ring. Then, a central limit theorem is proved that enables us to construct asymptotic Gauss tests of the vector of Minkowski functionals.

Finally, efficient computer algorithms implementing the above method are touched upon. Simulation results are presented illustrating the accuracy of the method.

Point-stationarity

Hermann Thorisson

Let N^o be the Palm version of a stationary Poisson process N in R^d, that is, N^o has the same distribution as $N + \delta_0$. Consider the following problem: when $d > 1$, is there some non-randomized way of shifting the origin of N^o from the point at the origin to another point T so that the distribution of N^o does not change?

This is clearly possible when $d = 1$, since then the intervals between points are i.i.d. exponential and remain so when the origin is shifted to the nth point on the right (or on the left) of the point at the origin. But what about $d > 1$?

The answer is yes. According to [1] we only need to find a bijective point-shift. But this is not so easy. Olle Häggström provided the first example of such a point-shift in 1999: shift to the closest point if that point has the point at the origin as its closest point, otherwise stay where you are. But is there a strictly non-zero point-shift?

In [2] it is shown that if the points of the process can be joined in a shift-invariant way into a single one-ended tree, then shifting along the succession line would yield a class of non-zero bijective point-shifts. Moreover, in [2] a tree-construction is given which yields a single tree when d = 2 and 3, but the construction yields only a forest when $d > 3$. In [3] a tree-construction is given which yields a single tree in any dimensions for the Poisson process, and [4] does the same for processes with isometry-invariant distributions.

If we go further beyond the Poisson case, a more general problem concerns the concept of ,,point-stationarity". Intuitively, point-stationarity means that the behaviour of a point process N^o relative to a given point of the process is independent of the point selected as origin. Formally, this concept is defined in [1] to be distributional invariance under bijective point-shifts ,,against any independent stationary background" and shown to be the characterizing property of the Palm version N^o of any stationary point process N in R^d. A natural question is whether the definition of ,,point-stationarity" can be reduced to distributional invariance under non-randomized bijective point-shifts. An approach to this problem will be outlined.

References

[1] Thorisson, H.: Coupling, Stationarity, and Regeneration. Springer, NY, (2000).

[2] Ferrari, P.A., Landim, C., and Thorisson, H.: Poisson trees, succession lines and coalescing random walks. Ann. Ist. H. Poincare, Prob. et Statist. (to appear).

[3] Holoroyd, A.E. and Peres, Y.: Trees and matchings from point processes. Elect. Comm. In Probab., 8:17-27, (2003).

[4] Timar, A.: Tree and grid factors of general point processes. Preprint.

Some (challenging) open problems in queueing and Markovian control

Henk Tijms

In queueing applications one has sometimes to compute the state probabilities of the finite-capacity model for several values of the finite buffer. In fact one faces the computation of the equilibrium distribution of a natural truncation of a denumerable Markov chain with state space $0, 1, \ldots$. This raises the question under which conditions there is some integer K with $0 < K < N$ such that the first K state probabilities of the finite-state truncated Markov chain with state space $0, 1, \ldots, N$ are proportional to the corresponding state probabilities in the infinite-state Markov chain. This property holds amongst others for the $M/G/1/s+C$ queue with $s = 1$ and finite waiting room C and the $M/M/s/s+C$ queue with $s > 1$ and finite waiting room C when K is taken equal to $s + C - 1$. Also this property holds for some queues with batch-Poisson input when the batch size satisfies some regularity conditions.

Another interesting problem arises in controlled queueing systems. For example, consider a Erlang-loss type of queueing model in which two types of customers arrive according to independent Poisson processes. Any customer finding upon arrival no free server is lost. In case one wishes to minimize a weighted sum of the loss rates of the two customers types, one can use Markovian control to find an optimal rule when the service times are exponential. It appears that the optimal rule is of the L_i type: accept a customer of type 1 (say) always when there is a free server and accept a customer of type 2 finding upon arrival i customers of priority type 1 present only if a server is free and less than L_i customers of type 2 are present. Simulation experiments reveal that the weighted loss rate of this optimal control rule is nearly-insensitive to the form of the underlying service-time distributions. This finding raises interesting theoretical questions.

Edited by Volker Schmidt

Participants

Prof. Dr. Venkat Anantharam
ananth@eecs.berkeley.edu
Department of Electrical
Engineering and Computer Sciences
University of California
Berkeley, CA 94720 – USA

Prof. Dr. Soren Asmussen
asmus@imf.au.dk
Department of Theoretical Statist.
Institute of Mathematics
University of Aarhus
Ny Munkegade
DK-8000 C Aarhus

Dr. Francois Baccelli
baccelli@sophia.inria.fr
francois.baccelli@ens.fr
Laboratoire de Mathematiques
de l'Ecole Normale Superieure
U.R.A. 762
45, rue d'Ulm
F-75005 Paris

Prof. Dr. Nicole Bäuerle
baeuerle@stochastik.uni-hannover.de
Institut für Mathematische Stochastik
Universität Hannover
Welfengarten 1
D–30167 Hannover

Dr. Bartek Blaszczyszyn
blaszczy@dmi.ens.fr
Ecole Normale Superieure
DI
45 rue d'Ulm
F-75230 Paris Cedex 05

Prof. Dr. Alexander Borovkov
borovkov@math.nsc.ru
Institute of Mathematics SO RAN
Koptjug pr. 4
Novosibirsk, 630090 – Russia

Prof. Dr. Maury Bramson
bramson@math.umn.edu
School of Mathematics
University of Minnesota
127 Vincent Hall
206 Church Street S. E.
Minneapolis, MN 55455 – USA

Prof. Dr. Hans Daduna
daduna@math.uni-hamburg.de
Fachbereich Mathematik - ST
Universität Hamburg
Bundesstr. 55
D–20146 Hamburg

Prof. Dr. Daryl J. Daley
daryl@maths.anu.edu.au
Centre for Mathematics & its Applications
Australian National University
G.P.O. Box 4
Canberra ACT 0200 – Australia

Prof. Dr. Paul Dupuis
dupuis@dam.brown.edu
Division of Applied Mathematics
Brown University
Box F
182 George Str.
Providence, RI 02912 – USA

Prof. Dr. Serguei Foss
S.Foss@ma.hw.ac.uk
Department of Actuarial Mathematics
and Statistics
Heriot-Watt University
Riccarton
GB-Edinburgh EH14 4AS

Dr. Andreas Frey
a.frey@siemens.com
Siemens AG
Entwicklungszentrum Ulm
ICM N PG U SE A1
D–89025 Ulm

Prof. Dr. Bruno Gaujal
Bruno.Gaujal@ens-lyon.fr
ENS Lyon
LIP
46, Allee d'Italie
F-69007 Lyon Cedex

Prof. Dr. Peter Glynn
glynn@leland.stanford.edu
glynn@stanford.edu
Management Science and Engineering
Terman Engineering Center
Room 313
Stanford University
Stanford CA 94305-4026 – USA

Prof. Dr. Olle Häggström
olleh@math.chalmers.se
Department of Mathemtics
Chalmers University of Technology
S-412 96 Goteborg

Prof. Dr. Lothar Heinrich
heinrich@math.uni-augsburg.de
Institut für Mathematik
Universität Augsburg
D–86159 Augsburg

Prof. Dr. Arie Hordijk
hordijk@math.leidenuniv.nl
Mathematisch Instituut
Universiteit Leiden
Postbus 9512
NL-2300 RA Leiden

Prof. Dr. Mark Huber
mhuber@math.duke.edu
Department of Mathematics and the
Institute of Statistics and
Decision Sciences, Duke University
228 Physics
Durham NC 27708-0320 – USA

Prof. Dr. Wilfrid S. Kendall
w.s.kendall@warwick.ac.uk
Department of Statistics
University of Warwick
GB-Coventry CV4 7AL

Prof. Dr. Takis Konstantopoulos
takis@ece.utexas.edu
tk@math.upatras.gr
Department of Mathematics
University of Patras
26500 Patras – Greece

Prof. Dr. Ger Koole
koole@cs.vu.nl
Department of Mathematics
Vrije University
De Boelelaan 1081 a
NL-1081 HV Amsterdam

Dr. Günter Last
g.last@math.uni-karlsruhe.de
Institut für Math. Stochastik
Universität Karlsruhe
D–76128 Karlsruhe

Prof. Dr. Jean Mairesse
mairesse@liafa.jussieu.fr
LIAFA
Universite Paris VII
Case 7014
2 place Jussieu
F-75251 Paris Cedex 05

Prof. Dr. Armand Makowski
armand@eng.umd.edu
armand@isr.umd.edu
Department of Electrical and
Computer Engineering
College of Engineering
University of Maryland
College Park MD 20742 – USA

Dr. James Martin
James.Martin@liafa.jussieu.fr
LIAFA
Universite Paris VII
Case 7014
2 place Jussieu
F-75251 Paris Cedex 05

Prof. Dr. David McDonald
dmdsg@omid.mathstat.uottawa.ca
Department of Mathematics
University of Ottawa
585 King Edward Ave., POB 450
Ottawa, Ont. K1N 6N5 – Canada

Prof. Dr. Sean Meyn
meyn@vivc.edu
Coordinated Science Laboratory
University of Illinois
1308 W. Main Street
Urbana, IL 61801 – USA

Prof. Dr. Ilkka Norros
Ilkka.Norros@vtt.fi
VTT Information Technology
Tietoliikenne
P.O.Box 1202
FIN-02044 VTT

Dipl.Math. oec. Ursa Pantle
pantle@mathematik.uni-ulm.de
Abteilung Stochastik
Universität Ulm
Helmholtzstr. 18
D–89081 Ulm

Prof. Dr. Philip K. Pollett
pkp@maths.uq.edu.au
Department of Mathematics
University of Queensland
Brisbane, Queensland 4072 – Australia

Prof. Dr. Balaji Prabhakar
balaji@stanford.edu
Dept. of Electrical Engineering and
Computer Science
Packard Building, Room 269
Stanford University
Stanford CA 94305-9510 – USA

Prof. Dr. Sidney I. Resnick
sid@orie.cornell.edu
sir1@cornell.edu
ORIE
Cornell University
Rhodes Hall 214
Ithaca, NY 14853 – USA

Prof. Dr. Ad Ridder
aridder@feweb.vu.nl
Department of Econometrics
Vrije University
De Boelelaan 1105
NL-1081 HV Amsterdam

Prof. Dr. Rhonda Righter
RRighter@ieor.berkeley.edu
Department of Industrial
Engineering and Operations Research
University of California
Etcheverry Hall
Berkeley, CA 94720-1777 – USA

Prof. Dr. Tomasz Rolski
rolski@math.uni.wroc.pl
Institute of Mathematics
Wroclaw University
pl. Grunwaldzki 2/4
50-384 Wroclaw – Poland

Prof. Dr. Manfred Schäl
schael@uni-bonn.de
Institut für Angewandte Mathematik
Universität Bonn
Wegelerstr. 6
D–53115 Bonn

Dr. Martin Schlather
Martin.Schlather@uni-bayreuth.de
schlather@cu.lu
Laboratoire de Mathematique
Universite du Luxembourg
162 A, avenue de la Faiencerie
L-1511 Luxembourg

Prof. Dr. Hanspeter Schmidli
schmidli@math.ku.dk
Laboratory of Actuarial Mathematics
University of Copenhagen
Universitetsparken 5
DK-2100 Copenhagen

Dipl.-Math.oec. Hendrik Schmidt
hendrik@mathematik.uni-ulm.de
Abteilung Stochastik
Universität Ulm
Helmholtzstr. 18
D–89081 Ulm

Prof. Dr. Volker Schmidt
volker.schmidt@mathematik.uni-ulm.de
Abteilung Stochastik
Universität Ulm
Helmholtzstr. 18
D–89081 Ulm

Prof. Dr. Richard F. Serfozo
rserfozo@isye.gatech.edu
richard.serfozo@isye.gatech.edu
School of Industrial Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205 – USA

Prof. Dr. Adam Shwartz
adam@ee.technion.ac.il
Dept. of Electrical Engineering
TECHNION
Israel Institute of Technology
Haifa 32000 – ISRAEL

Dr. Flora M. Spieksma
spieksma@math.leiden.univ.nl
Department of Mathematics and
Computer Science
Rijksuniversiteit Leiden
Postbus 9512
NL-2300 RA Leiden

Prof. Dr. Evgueni Spodarev
spodarev@mathematik.uni-ulm.de
Abteilung Stochastik
Universität Ulm
Helmholtzstr. 18
D–89081 Ulm

Prof. Dr. Hermann Thorisson
hermann@hi.is
Science Institute
University of Iceland
3 Dunhaga
IS-107 Reykjavik

Prof. Dr. Henk C. Tijms
tijms@econ.vu.nl
tijms@feweb.vu.nl
Department of Econometrics
Vrije University
De Boelelaan 1105
NL-1081 HV Amsterdam