KEYNOTE SPEAKERS

Gareth Roberts
Some recent developments in scaling of Metropolis-Hastings Algorithms

This talk will survey the area of optimal scaling of MCMC algorithms, focusing largely on Metropolis and Langevin methods. The problem will be studied by the optimisation of infinite dimensional limits of algorithms, which turn out (usually) to be diffusion processes. In many circumstances, optimally scaled algorithms can be characterised as the algorithm achieving particular overall acceptance probabilities. For instance, the optimal acceptance rate of Metropolis methods is often 0.234. The talk will discuss when and why such criteria hold, and describe some interesting cases where such target acceptance rate criteria break down.

Harvard Rue
Alternatives to MCMC based inference for spatial models

MCMC based inference for spatial models is not without practical problems due to slow convergence and/or huge computational requirement. In this talk I will discuss strategies to circumvent these problems by doing approximate inference for spatial models: spatial GLM's and hierarchical GMRF models in particular. The aim is to approximate posterior marginals of the parameters of interest. Our approach is to use nested Laplace-approximations, and has the advantage of being fast to compute and can be computed nearly instant compared to any MCMC alternative.
The accuracy is surprisingly good and the error is relative, not additive as for any MCMC approach.

Eric Moulines
ODE methods for Markov chain stability with applications to MCMC
(joint work with G. Fort, S. Meyn, P. Priouret)

Fluid limit techniques have become a central tool to analyze queueing networks over the last decade, with applications to performance analysis, simulation, and optimization. In this paper some of these techniques are extended to a general class of skip-free Markov chains. As in the case of queueing models, a fluid approximation is obtained by scaling time, space, and the initial condition by a large constant. The resulting fluid limit is the solution of an ODE in "most" of the state space. Stability and finer ergodic properties for the stochastic model then follow from stability of the set of fluid limits. Moreover, similar to the queueing context where fluid models are routinely used to design control policies, the structure of the limiting ODE in this general setting provides an understanding of the dynamics of the Markov chain "at infinity". These results are illustrated through application to Markov Chain Monte Carlo.

Christophe Andrieu
On the pseudo-marginal Hastings-Metropolis algorithm (joint with G. O. Roberts)

In numerous scientific areas where the probabilistic formalism is used, it is often the case that probabilities of interest do not admit closed-form expressions, but can be thought of as being marginal distributions of probabilities which admit simple closed form expressions, sometimes up to a normalising constant. In a Bayesian framework for example, it is often the case that the expression for the posterior distribution of the parameter/variable of interest does not admit an analytical expression, whereas the joint posterior distribution of the parameter of interest and a set of adequately chosen latent variables admits a simple analytical expression. The use of Monte Carlo methods, and in particular MCMC algorithms, allow one, at least in principle, to easily tackle the additional marginalising problem introduced by the presence of the latent variables. However this is often at the expense of a loss of efficiency.
In the context of the Hastings-Metropolis algorithm, the pseudo-marginal approach allows for this loss of efficiency to be alleviated by introducing some form of `local adaptation' that can lead to efficient algorithms, while requiring little tuning or expertise. A theoretical analysis of some of these algorithms, which can be shown to perform well under very mild assumptions, will be presented as well as numerical simulations in various contexts, including in particular reversible jump MCMC or trans-dimensional MCMC algorithms.

Samuel Kou
Equi-energy sampler: From statistical inference to statistical mechanics to protein folding

We introduce a new sampling algorithm, the equi-energy sampler, for efficient statistical sampling and estimation. Complementary to the widely used temperature-domain methods, the equi-energy sampler, utilizing the temperature-energy duality, targets the energy directly. The focus on the energy function not only facilitates efficient sampling, but also provides an efficient means for statistical estimation, for example, the calculation of the density of states and microcanonical averages in statistical mechanics. The equi-energy sampler is applied to a variety of problems, including motif sampling in computational biology and protein folding in biophysics. This work is joint with Qing Zhou and Wing Wong.

INVITED SESSION 1

Jeffrey Rosenthal
Adaptive MCMC: A Java Applet's Perspective

When sampling from complicated probability distributions, a wide variety of MCMC algorithms are available. It is tempting to have the computer automatically "adapt" the algorithm while it runs, to improve and tune on the fly. However, natural-seeming adaptive schemes often fail to preserve the stationary distribution, thus destroying the fundamental ergodicity properties necessary for MCMC to be valid. In this talk, we review adaptive MCMC, and present simple conditions which ensure ergodicity (proved using intuitive coupling constructions, jointly with G.O. Roberts). The ideas are illustrated using the very simple example animated by the java applet at: probability.ca/jeff/java/adapt.html.

Mylene Bedard
Optimal acceptance rates for Metropolis algorithms: moving beyond 0.234

In this talk, we shall optimize the efficiency of random walk Metropolis algorithms for multidimensional target distributions with scaling terms possibly depending on the dimension. We show that when there does not exist any component having a scaling term significantly smaller than the others, the asymptotically optimal acceptance rate is the well-known 0.234. We also show that when this condition is not met the limiting process of the algorithm is altered, yielding an asymptotically optimal acceptance rate which might drastically differ from the usual 0.234. In particular, we prove that as the d increases the sequence of stochastic processes formed by say the component of each Markov chain usually converges to a Langevin diffusion process with a distinct speed measure, except in particular cases where it converges to a one-dimensional Metropolis-Hastings algorithm with a singular acceptance rule. We also discuss the use of inhomogeneous proposals, which might reveal essential in specific cases.

Ajay Jasra
Non-Linear Markov chain Monte Carlo via Self Interacting Approximations

In his talk we introduce the idea of non-linear Markov chain Monte carlo (MCMC) to solve the problem of simulating from a probability measure π. Non-linear Markov kernels (e.g. Del Moral (2004); Del Moral & Doucet (2003)) can be constructed to admit π as an invariant distribution and have superior mixing properties to ordinary (linear) MCMC kernels. However, such non-linear kernels cannot be simulated exactly, so, in the spirit of particle approximations of Feynman-Kac formulae (Del Moral 2004), we construct approximations of the non-linear kernels via self interacting Markov chains (Del Moral & Miclo 2004; 2006) (SIMC). We present several non-linear kernels and demonstrate that, under verifiable conditions, the self-interacting approximations can be simulated so that they exhibit a strong law of large numbers; our proof technique is via the Poisson equation and Foster-Lyapunov conditions. We investigate the performance of our approximations with some simulations, combining the methodology with population-based Markov chain Monte Carlo. This is joint work with Christophe Andrieu (Bristol), Arnaud Doucet (UBC) and Pierre Del Moral (Nice).

INVITED SESSION 2

Nial Friel
Recursive computing and simulation-free inferencefor Markov random fields

In this talk we present recursive algorithms which allows a variety of tasks to be performed exactly for Markov random fields (MRFs) including, sampling of MRFs, computing marginal distributions, maximising MRFs and computing marginal likelihoods. All of these can be carried out exactly for lattices where the number of rows or columns of the lattice is no greater than 20. In addition, we present approximate methods to extend these exact results to larger lattice sizes. We present results for a variety of datasets. This work is joint with Havard Rue (NTNU, Trondheim).

Merrilee Hurn
MCMC for estimating galaxy redshift

This talk describes an application in astronomy where the goal is to estimate galaxy redshift based on a comparatively cheap set of data (photometric data) calibrated by more expensive data (spectroscopic data) for a small set of galaxies. One of the problems which arises is the occasional catastrophic estimation error, often with a deceptively small associated credible interval, and poor mixing may be to blame for this, at least in part. (Joint work with Peter Green and Fahima Al-Awadhi).

Mark Huber
Perfect simulation from continuous state spaces with the Randomness Recycler

The Randomness Recycler (RR) protocol has been utilized to create perfect sampling algorithms for several problems on discrete state spaces. In this talk I will show how to examine the RR methodology to arbitrary spaces. As an example, this new approach is applied to sampling from the autonormal model, which arises in Bayesian image analysis. Under certain ranges of parameters of the problems this yields an interruptible linear time algorithm for generating random variates.

Wilfrid Kendall
Perfect simulation: a survey

This talk will survey a number of perfect simulation techniques, updating the survey in Kendall (2005), and then describe recent work on general perfect simulation joint with Connor.
Reference:
Kendall, W.S. (2005) "Notes on Perfect Simulation", pp 93-146, *Markov chain Monte Carlo: Innovations and Applications* edited by Kendall, Liang and Wang, World Scientific, Singapore.

INVITED SESSION 3

Gersende Fort
Criteria for subgeometric ergodicity of strong Markov processes (jointly with G.O. Roberts, R. Douc, A. Guillin)

Criteria for subgeometric ergodicity of strong Markov processes (jointly with G.O. Roberts, R. Douc, A. Guillin). We introduce conditions for *f*-ergodicity of strong Markov processes at a subgeometric rate. The first condition is couched in terms of delayed return-time to a petite set. Nevertheless, this condition is not designed for practical applications. Hence, we also derive conditions expressed in terms of inequality on the extended generator, and is analogous to the so-called drift inequality in the discrete-time case. This condition is easy to check in many applications. Applications to specific processes will be considered. We will first study stochastic differential equations with drift vector b verifying

$$< b(x), x/|x| > \leq -r|x|^p$$

for large $|x|$ and some $0 < p < 1$. We will then consider Langevin tempered diffusion in which the diffusion matrix is a scalar matrix with coefficient π^{-2d}, $d > 0$, where π is the invariant distribution of the diffusion. We prove that even when the target distribution is heavy tailed, a convenient choice of the temperature d involves geometric ergodicity of the process.

Peter Green
Branching process Monte Carlo

This talk is an exploration of the possible role for branching processes and related models in Monte Carlo simulation from a complex distribution, such as a Bayesian posterior. The motivation is that branching processes can support antithetic behaviour in a natural way by making offspring negatively correlated, and also that branching paths may assist in navigating past slowly-mixing parts of the state space. The basic theory of branching processes as used for sampling is established, including the appropriate analogue of global balance with respect to the target distribution, evaluation of moments, in particular asymptotic variances, and a start on the spectral theory. Although our model is a kind of 'population Monte Carlo', it should be noted that it has virtually nothing to do with particle filters, etc. Our target is not sequentially evolving, and we rely on ergodicity for convergence of functionals of the target distribution, rather than using importance sampling. This is joint work with Antonietta Mira (University of Insubria, Varese, Italy).

Omiros Papaspiliopoulos
Stability of the Gibbs sampler of Bayesian hierarchical models

In this talk we characterise the convergence of the Gibbs sampler which samples from the joint posterior distribution of parameters and missing data in hierarchical linear models with arbitrary symmetric error distributions. We show that the convergence can be uniform, geometric or sub-geometric depending on conditions on the relative tails of the prior distribution and the likelihood. We show that these conditions also determine the robustness properties of the model in presence of large likelihood-prior conflict. We show how to obtain much more stable algorithms by re-parametrisations and we apply our theory to characterise the convergence of the Gibbs sampler on latent Gaussian process models.

INVITED SESSION 4

Andrew Stuart
Sampling Conditioned Diffusions

There are a wide variety of applications which can be cast as sampling problems for conditioned SDEs (diffusion processes). Examples include nonlinear filtering in signal processing, data assimilation in the ocean/atmosphere sciences, data interpolation in econometrics, and finding transition pathways in molecular systems. In all these examples the object to sample is a path in time, and is hence infinite dimensional. We describe an abstract MCMC method for sampling such problems, based on generalizing Metropolis adjusted Langevin algorithms to infinite dimensions. This leads naturally to the study of stochastic reaction-diffusion equations which, in their invariant measure, sample from the required distribution. Furthermore, the study of preconditioning in this context leads to some interesting new infinite dimensional semilinear evolution equations. We give an overview of the subject area, describing the analytical and computational challenges, and illustrating applicability of the techniques being developed.

Jochen Voss
An MCMC Method for Sampling Diffusion Bridges

We present an MCMC Langevin algorithm for sampling nonlinear diffusion bridges. The method is based on the idea to derive a stochastic partial differential equation (SPDE) on the space of all bridge paths which has the target bridge as its stationary distribution. The SPDE is found by transfering the finite dimensional Langevin method to this setting. We discretise the SPDE in the algorithmic time direction using an implicit scheme, parameterised by a parameter $0 \leq \theta \leq 1$, and correct for the bias with a Metropolis-Hastings accept/reject mechanism. In practice, the algorithm performs well only when $\theta = ½$ for other values the acceptance rate is negligible. Numerical results illustrating this phenomenon are presented, and a theoretical justification is given. The key issue is producing proposals with the correct quadratic variation. Joint work with Alexandros Beskos, Gareth Roberts, Andrew Stuart.

INVITED SESSION 5

David Stephens
MCMC for Levy Process Models of Stochastic Volatility

MCMC algorithms for Levy process driven stochastic volatility models are now well established for the case where the latent volatility model is driven a pure jumps process, and the observation model is driven by a standard Brownian motion. We review extensions to the most commonly implemented "Barndorff-Nielsen and Shephard" models, including observation processes that are themselves Levy processes with jump components, and are well established in the mathematical finance literature. We illustrate the fitting of competing models to a standard test data set, the S&P 500 index, and examine implications for option pricing. (This is joint work with Matthew Gander and Wing Yip, Imperial College London).

Chris Holmes
MCMC methods for inferring ancestry within subpopulations undergoing genetic transfer

We discuss our recent work in developing MCMC methods to infer phylogenies under a coalescent model of evolution which accommodates genetic transfer (migration) between known subpopulations. Our motivating example is in the analysis of HIV-1 data collected from homosexual men and injecting drug users. Of particular scientific interest is the extent of genetic transfer between the two viral subpopulations and the time to most recent common ancestor. We shall discuss the complexities in modelling this process, the simplifying assumptions that are currently necessary, and the direction of future work in this field.

Sylvia Richardson
(Joint work with M. Zucknick and C. Holmes)
MCMC methods for Bayesian variable selection in cases where structured dependence among the covariates is present: application to gene expression data

In many genomic applications, data sets with many more variables than samples are common. One example are gene expression microarray studies where the number of genes p available is usually many thousand whereas the number of observations n is much smaller. In this context variable selection is an important problem, since there is much interest in finding parsimonious regression models that include only a small set of genes which can be interpreted in biological terms (Sha et al 2004). Furthermore, the set of genes typically exhibit complex dependence structure, with groups of genes involved in the same pathway expected to have correlated expression. Bayesian variable selection models implement covariate uncertainty by a binary covariate indicator variable γ which determines which variables are selected. Because we are in a "large p, small n" situation, the posterior distribution over the model space of variable dimensions is multi-modal. Also, full posterior inference for the entire model space of size 2^p is not feasible if p is larger than about 20. Hence, Markov chain Monte Carlo methods are rather used as stochastic search algorithms with the aim to quickly find many regions of high posterior probability.

For these reasons, the Markov chain needs to move quickly around the support of the posterior distribution. We will focus on the logistic regression set-up with binary outcome and a large number of regressors. Following the formulation discussed by Holmes and Held (2006), a joint update of the covariate indicator γ and of the regression coefficients is implemented. We discuss different strategies for updating the covariate indicator, which exploit the dependence structure among the covariates. The mixing and convergence performances of the resulting Markov chains are evaluated and compared to standard samplers which only update one or two variables at a time in both a simulation study and in an application to a real gene expression data set. References:

C.C. Holmes and L. Held. Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Analysis*, 1:145-168 (2006)

Sha, N., et al. Bayesian variable selection in multinomial probit models to identify molecular signatures of disease stage. *Biometrics*, 60:812-819 (2004).

Eleisa Heron

Parameter estimation for a stochastic model of a gene regulatory network

Several systems in biology can be described by a negative feedback regulatory mechanism. For example, the circadian rhythm in Arabidopsis thaliana results from the dynamic process of regulation of gene expression. A stochastic model for such a regulatory network, consisting of coupled differential equations, is introduced. The parameters of the model have meaningful biological interpretations and the aim is to estimate the parameters of the model using various MCMC techniques.

POSTERS

Teresa Barata
Photo-identification of bottlenose dolphins using MCMC

In animal behaviour and ecology the ability to recognize individuals is very important and most long-living animals can be identified by natural marks. In the case of marine mammals photo-identification must be used, and for bottlenose dolphins recognition is done via their dorsal fins. Here we propose to identify individual bottlenose dolphins by means of a parametric model for their dorsal fin shape. Our model consists of shape parameters (to account for the distance and (3-D) angles to the dolphin being photographed) yielding quite a complicated curve. The raw data consists of photographs of dorsal fins from which the outline must be extracted. The pixel coordinates in this outline will be the datapoints to which the model must be fitted. To do this we use a Bayesian regression approach to model fitting, with errors in both the x and y variables. In a previous approach we considered using bivariate normal errors, whereas here we propose a more realistic error structure, that gives rise to square, instead of circular contour plots. Extra consideration must be given while doing estimation and in this talk we demonstrate how MCMC methods can be used to fit the model to data.

Andrew Golightly
Bayesian inference for nonlinear diffusion models observed with error

Diffusion processes governed by stochastic differential equations (SDEs) are a well established tool for modelling continuous time data from a wide range of areas. Consequently, techniques have been developed to estimate diffusion parameters from partial and discrete observations. Unfortunately, likelihood based inference can be problematic as closed form transition densities are rarely available. One widely used solution adopts the treatment of Pedersen (1995) and involves the introduction of latent data points between every pair of observations to allow an Euler-Maruyama approximation of the true transition densities to become accurate. In recent literature, Markov chain Monte Carlo (MCMC) methods (see for example, Golightly and Wilkinson (2005)) have been used to sample the posterior distribution of latent data and model parameters; however, naive schemes suffer from a mixing problem, highlighted by Roberts and Sramer (2001), that worsens with the degree of augmentation. We therefore use the sequential MCMC scheme of Golightly and Wilkinson (2006) which relies on a joint update of parameters and latent values and whose performance is not adversely affected by the amount of augmentation. We illustrate the methodology by estimating parameters governing an auto-regulatory gene network, using partial and discrete data that is subject to measurement error.

Maria Kalli
Mixtures of Dirichlet Processes (MDP) and the Slice Sampler

Mixtures of Dirichlet Processes (MDP) have been widely used as a method of overcoming the discreteness of the Dirichlet Process (DP). The two approaches taken to sample from the Dirichlet measure are: the marginal approach (Escobar and West 1995) where the measure is integrated out within the Gibbs sampler via a clever use of the Polya Urn construction of the DP and the conditional approach (see Ishwaran and Zarepour 2000, 2002) which makes use of the infinite sum construction of the DP (see Sethuraman 1994). The ways around this infinite sum construction are either approximations or truncations (Ishwaran and Zarepour 2000) or using the retrospective sampler (see Papaspiliopoulos and Roberts 2005). The retrospective sampler deals with the infinite sum directly, via use of reversible jump steps. We introduce a simpler sampler, which instead of using reversible jumps, introduces an auxiliary variable and incorporates the slice sampler within the construction of the posteriors for the Gibbs sampler (see P. Damien, J. Wakefield, S.G.Walker 1999). The new algorithm works with the infinite sum construction of the DP from the very beginning and by introducing auxiliary variables the Gibbs sampler updating is done within finite sets.

Michaelis Kolossiatis
Bayesian nonparametric modelling of spatial data

An area of Bayesian Statistics that has attracted much attention recently is non- or semi-parametric modelling. In many of those applications, a non-parametric prior process is expressed by a Dirichlet Process (and more recently a Normalised Inverse-Gaussian Process) for an unknown distribution. These models seem to be the natural choice in a hierarchical model. Finally, these models take great advantage of the computational benefits of MCMC methodology. Taking the Dirichlet Process as an example, we can see that a way of constructing a random probability measure is by normalising a random measure (the Dirichlet Process can be seen as a normalised Gamma Process). Also, the Normalised Inverse-Gaussian Process is a normalised inverse-Gaussian Process. The aim of my project will be to exploit the infinite divisibility of the underlying random measure (in our examples, the Gamma and the Inverse-Gaussian Processes, for which the infinite divisibility holds), in order to construct random probability measures that are identically distributed, but not independent. I will study the properties of those models, and in the future, try to apply them in spatial modelling.

Theodore Kypraios
Robust MCMC algorithms for a stochastic epidemic model & a fully Bayesian Analysis of the 2001 Foot-And-Mouth Epidemic occurred in the UK

In general, inference problems for disease outbreak data are complicated by the facts that (i) the data are inherently dependent and (ii) the data are usually incomplete in the sense that the actual process of infection is not observed. We adopt a Bayesian approach and apply Markov Chain Monte Carlo (MCMC) methods in order to make inference for the parameters of interest (infection and removal rates). Once the size of the data set increases, the standard methods perform poorly. Therefore, apart from centered reparameterisation we extend the Non-Centered and partially Non-Centered algorithms presented in Neal and Roberts (2005). Finally, we adopt a fully Bayesian approach to analyze the Foot-and-Mouth disease occurred in 2001 in the UK.

Demetris Lamnisos
Bayesian Variable selection in classification problems

We are interested in the variable selection problem for the binary probit model in cases where the number of predictor variables p exceeds the sample size n. Data in which the number of predictors variables far exceeds the sample size arise in DNA Microarrays expression studies. All the possible subsets of the predictor variables determine the model space and each subset can be parameterized via a one to one mapping to a binary vector. In the case of the probit model the model specific parameters are the regression coefficients. Our aim is to construct efficient MCMC algorithms with limiting distribution the joint posterior distribution of the model space and the model parameters. For this reason, we apply both the higher order and conditional maximization method discussed by Brooks et al. (2003) in order to propose, in each iteration, the vector of regression coefficients for the new model. These methods automatically determine the location and scaling for the proposal distribution. In this case the proposal parameters are functions of the current state of the MCMC. The mixing and convergence performance of these algorithms are evaluated and compared to a sampler discussed by Holmes and Held (2006) in an application to a real gene expression data set. The last algorithm updates the regression coefficients from its full conditional distribution given that a change to the predictors variables has been accepted.

Beatriz Penaloza
The Uncertainty of Binary Data Models

Modelling binary data implies the choice of link function and the choice of covariate set. The covariate set uncertainty has been explored in the literature, mainly through Bayesian Model Averaging (BMA). One of the main algorithms used in this context is the Reversible Jump Monte Carlo Markov Chain (RJMCMC) proposed by Green (1995). In the case of link function uncertainty less research has been done. We extend the Albert and Chib (1993) proposal and define the shape of the binary data model in terms of a mixing parameter λ, which let the data determine its most likely distribution. We propose an algorithm that considers both uncertainties within a Bayesian framework. The regression model for binary data is defined as a scale mixture of normals linear model based on latent continuous data to deal with link function uncertainty. This approach allows for flexible shapes which do not necessarily correspond to a known model. The set of covariates that are used in the model is also left free and estimated from the data using the RJMCMC algorithm. We present the extended algorithm. The performance of the model and some interesting results are illustrated in an applied exercise. (This is joint work with M.J. Steel).
References:
(1) Albert, J.H., and Chib, S.: Bayesian analysis of binary data and polychotomous response data. *Journal of the Americal Statistical Association 88,*422 (1993) , 669-679.
(2) Green, P.J.: Reversible jump Markov Chain Monte Carlo computation and Bayesian model determination. *Biometrika 82,*4 (1995), 711-732.

Volker Schmid
A non-parametric approach for PK models in DCE-MRI

Dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) has become an important source of information to aid cancer diagnosis. First approaches analysed semi-quantitative statistics like area under the signal curve, onset time etc. Recent approaches fit non-linear models, convoluted with the arterial input function (AIF), to the contrast agent concentration time series, either in each pixel or for a whole region of interest. These methods provide meaningful biological parameters, but have issues in convergence, (de-)convolution of the AIF and in properly fitting observed data. To overcome this problems, we use a B-spline smoothing approach to model the data in a non-parametric way. To incorporate the AIF, we convolute the AIF with the B-spline design matrix. Modelling the resulting deconvoluted flow, biological parameters can be achieved like with the parametric approaches. We present results on a study on breast cancer treatment.

Chris Sherlock
Optimal scaling for the Metropolis-Hastings random-walk on unimodal elliptically symmetric targets (with Gareth Roberts)

A major problem in the implementation of Metropolis algorithms is the choice of scaling. Optimal scaling results based on acceptance rate criteria lead to generic implementation rules which are easy to apply in practice. For instance one widely adopted method attempts to tune the acceptance rate to 0.234. Existing theoretical results cover specific situations such as independent identically distributed targets, and are obtained through a limiting diffusion approximation to the random walk as dimension $d \to \infty$. Results to be detailed are based around exact closed forms for the expected acceptance rate and square jump distance, which apply whatever the dimension of the problem. They are obtained using a simpler approach and apply to unimodal spherically and elliptically symmetric targets. The method provides an intuition as to when the limiting acceptance rate of 0.234 is applicable, and to possible outcomes when it is not. In particular it can be shown that in many situations the limiting optimal acceptance rate is < 0.234. The method also yields a simple formula for limiting relative efficiencies between elliptical and spherical proposals on elliptical targets.

Tristan Marshall
Perfect Simulation for parameters of a partially observed diffusion process

Beskos et al. (2005 RSS B Read Paper) introduced a methodology for simulating a diffusion process and for performing parametric inference for a partially observed diffusion. In particular they demonstrated how to perform Bayesian inference for a parameter θ of such a diffusion; we can implement a form of `Data Augmentation' algorithm by alternately imputing the paths between observed points conditional on θ, and updating for θ conditional on these imputed paths. Here we look at adapting this data augmentation algorithm to achieve perfect simulation, so that we can draw from the exact posterior distribution of θ. The algorithm we implement is a form of Read-Once Coupling-From-the-Past; this approach avoids the need to either run the simulation in reverse or retain large amounts of path data in memory. It turns out that there is a very simple Perfect Simulation algorithm that will work for a wide class of diffusions, but its performance deteriorates rapidly with the size of the observed data. Future work will focus on finding algorithms that perform better but remain applicable to a reasonably large class of processes, we suggest some possible approaches.

Bruno Casella
Partially Implicit Langevin Schemes with MCMC Applications

Langevin algorithms are well known simulation methods for sampling from a given probability density π. The main idea is to construct a (Langevin) SDE which has an invariant measure with density π and perform some kind of discretization to mimic the dynamics of the SDE. The hope is that the invariant measure of the discretized chain is close to π, at least for sufficiently small discretization intervals Δ. This seems also a very natural way to construct efficient proposals for Metropolis Hasting algorithms. Surprisingly, it turns out that for certain choices of the π usual Euler discretization is not even ergodic (it is indeed transient), regardless the value of Δ we choose. This effect occurs in particular when π has lighter tails than Gaussian (light tails case). Clearly we can force the chain to be ergodic by adding a Metropolis-Hastings rejection step. However, in the light tails case, the resulting chain still fails to be geometrically ergodic. In order to overcome these problems, here we introduce a class of Langevin algorithms based on a more sophisticated discretization of the Langevin SDE which combines the Euler (explicit) approach and the implicit approach. We show that, by choosing appropriately the relative weight of the implicit component, we can generate a Markov chain which is geometrically ergodic also in the light tails case. Moreover geometric ergodicity is inherited by the Metropolis adjusted algorithm thus providing an efficient method to sample from light tails distributions.

Miguel Belmonte
Stochastic Conditional Duration model and Particle Filters

Duration models are a class of self-exciting models that aim at modelling time between financial transactions. A satisfactory explanation of inter-arrival time of transactions helps to understand the dynamics of stock prices. Furthermore, the intensity of transactions reflects the knowledge of traders about the market. Recently, a model in state space form, known as stochastic conditional duration (SCD), has been recently formulated in Bauwens and Veredas (2004). This new model allows sequential on-line inference about the hidden intensity by using particle filters. The particle filter is a simulation filter which approximates the filtering density once the parameters of a general state space model are known. Rather than assuming values of these parameters, we intend to estimate those by maximum likelihood. A smooth particle filter can be employed to approximate the likelihood function of sequential observations with depends on a vector of unknown parameters.