– 292 (revised) –

Perfect Simulation for the Area-Interaction Point Process

by

W.S. Kendall

April 1996

To appear in *Probability Perspective*
ed. by L. Accardi and C.C. Heyde, late 1997
World Scientific Press, Singapore.

Perfect Simulation for the Area-Interaction Point Process.

Wilfrid S. Kendall
Department of Statistics, University of Warwick
Coventry CV4 7AL, UK.

Abstract: Because so many random processes arising in stochastic geometry are quite intractable to analysis, simulation is an important part of the stochastic geometry toolkit. Typically, a Markov point process such as the area-interaction point process is simulated *(approximately)* as the long-run equilibrium distribution of a (usually reversible) Markov chain such as a spatial birth-and-death process. This is a useful method, but it can be very hard to be precise about the length of simulation required to ensure that the long-run approximation is good. The splendid idea of Propp & Wilson (1996) suggests a way forward: they propose a coupling method which delivers exact simulation of equilibrium distributions of (finite state-space) Markov chains. In this paper their idea is extended to deal with perfect simulation of attractive area-interaction point processes in bounded windows. A simple modification of the basic algorithm is described which provides perfect simulation of the repulsive case as well (which being non-monotonic might have been thought out of reach). Results from simulations using a `C` computer program are reported; these confirm the practicality of this approach in both attractive and repulsive cases. The paper concludes by mentioning other point processes which can be simulated perfectly in this way, and by speculating on useful future directions of research. Clearly workers in stochastic geometry should now wherever possible seek to incorporate the Propp and Wilson idea in their simulation algorithms.

KEYWORDS: PERFECT SIMULATION; SPATIAL BIRTH-AND-DEATH PROCESSES; AREA-INTERACTION POINT PROCESS; STRAUSS POINT PROCESS; EXCLUDED-AREA POINT PROCESS; POISSON POINT PROCESS; BOOLEAN MODEL; COUPLING; MARKOV CHAIN MONTE CARLO.

AMS 1991 SUBJECT CLASSIFICATION: 62M30, 60G55, 60K35.

1: Perfect Simulation.

Conventional Markov chain Monte Carlo techniques (MCMC) consist of long-run simulations of Markov chains: it is hoped that the long runs are sufficiently long that the Markov chain distribution is close to equilibrium. This vastly increases the range of applicability of Bayesian methods, as one can avoid computationally intractable issues relating to evaluation of the normalization constant of the posterior density. This is because simulating (approximately) the density as a long-run equilibrium distribution of a (typically reversible) Markov chain will be based ultimately on ratios of various posterior densities, thus avoiding reference to the normalization constant. For a discussion of these issues in a spatial context see Besag & Green (1993) (other papers in the same journal issue provide useful discussion of MCMC in other statistical contexts).

There is however a catch. How can one tell whether one is close enough to equilibrium? There are various empirical techniques and theoretical bounds, but the problem remains rather intransigent. This is a vexed question even in the simpler situation of simulating large-state-space Markov chains possessing substantial symmetry (arising in statistical mechanics and as discrete approximations to models in stochastic geometry).

A recent paper by Propp & Wilson (1996) shows how one can solve the problem completely in several cases, in which one can modify the simulation algorithm using coupling techniques so that the simulation itself provides a certificate stating when equilibrium has been achieved exactly! This allows *perfect simulation* of the equilibrium distribution of suitable finite-state-space Markov chains. This paper demonstrates the possibilities for stochastic geometry and spatial statistics, by showing how to exploit Propp and Wilson's ideas to develop an algorithm for perfect simulation of the equilibrium distribution of a certain spatial birth-and-death process, thus obtaining perfect simulation of the area-interaction point processes of Baddeley and Van Lieshout (1995).

The reader is referred to the Propp-Wilson paper for full details of their application of coupling techniques to simulation. Here we briefly discuss the simplest possible interesting case (though in fact the notation for the general case is identical).

Suppose that X is a two-state discrete-time Markov chain, with state-space $\mathcal{S}$ and stochastic matrix

$$P \quad = \quad \begin{pmatrix} 0 & 1 \\ 1/2 & 1/2 \end{pmatrix} .$$

We can view $\{X_t : t \geq 0\}$ as produced by a (discrete-time) *stochastic flow*: if $X_0 = i$ then $X_t = F_{0,t}(i)$, where the random maps $F_{s,t} : \mathcal{S} \to \mathcal{S}$ $(s \leq t)$ are defined by the composition $F_{s,t} = H_{s+1} \circ \ldots \circ H_t$, using an independent and identically distributed sequence $\{H_k : k \in \mathbb{Z}\}$ of transition maps $H_k : \mathcal{S} \to \mathcal{S}$ with distribution

$$\begin{aligned} \mathbb{P}[\, H_k(1) = 2 \,] \quad &= \quad 1 \,, \\ \mathbb{P}[\, H_k(2) = 1 \,] \quad &= \quad \mathbb{P}[\, H_k(2) = 2 \,] \quad = \quad 1/2 \,. \end{aligned}$$

Of course everything can be calculated exactly in this simple case, but consider how one might obtain the equilibrium distribution $(1/3,\, 2/3)$ of X by simulation. Orthodox simulation methods would approximate it by sampling $X_t(i) = F_{0,t}(i)$ for large t. Propp and Wilson point out that $Y_t(i) = F_{-t,0}(i)$ has the same distribution as $X(i)$ (though $\{Y_t : t \geq 0\}$ no longer forms a Markov chain) but also $Y_t(i)$ does not depend on i when $t \geq T_c$, for some (random, almost surely finite) T_c. Consequently $Y_\infty = \lim_{t\to\infty} Y_t$ exists and has the equilibrium distribution. Better yet, the random variable T_c can be defined as a *coalescence time* which can be viewed as a stopping time in a reverse filtration:

$$T_c \quad = \quad \inf\{t : F_{-t,0} \text{ has singleton image }\} .$$

Hence $F_{-T_c,0}(i) = Y_\infty$ gives an algorithm for sampling from the equilibrium distribution in finite (but random) time.

In our simple example $T_c = \inf\{t : H_{-t}(2) = 2\}$ has a Geometric distribution. More generally T_c is distributed as the first time to complete coupling of copies of X started at all possible initial positions and allowed to coalesce when they meet. This will work whenever the Markov chain is irreducible and ergodic. There is a relationship to work on estimation of the second eigenvalue using coupling techniques (see Chen's article in this volume, also Chen, 1994) as well as to other coupling ideas, as discussed in Propp and Wilson.

The essence of the above approach is as follows. When we can show that

(*a*) a Markov chain X_t on a state-space $\mathcal{S}$ is realized by a stochastic flow $F_{s,t} : \mathcal{S} \to \mathcal{S}$ $(s \leq t)$, in the sense that $X_t = F_{s,t}(X_s)$;

(*b*) there is *coalescence* of $F_{-T,0}$ for all sufficiently negative $-T$, in the sense that with probability one for all sufficiently negative $-T$ the result $F_{-T,0}(x)$ does not depend on the initial value $x \in \mathcal{S}$;

then an equilibrium distribution for the chain X exists and is given by

$$Y_\infty \quad = \quad \lim_{T\to\infty} F_{-T,0}(x) \,,$$

where the limit Y_∞ does not depend on the initial value x. We describe this situation as one where *perfect simulation* of the equilibrium distribution is possible, because in principle it then suffices to sample from $F_{-t,0}(x)$ for some $t > T$ (in which case $F_{-t,0}(x) = Y_\infty$).

For perfect simulation to be practical it must be computationally feasible to check whether or not coalescence has happened by (parsimonious!) inspection of the flow $F_{-T,-t}$ over a given simulation interval $(-t \in [-T, 0])$. Propp & Wilson (1996) observe that in case of monotonicity (as for example in the case of attractive interacting particle systems) this can be done at a cost which is comparable to conventional simulation techniques: if there are states $\mathbf{0}$ and $\mathbf{1}$ in $\mathcal{S}$ such that $F_{-T,0}(x)$ is sandwiched between $F_{-T,0}(\mathbf{0})$ and $F_{-T,0}(\mathbf{1})$ for all $x \in \mathcal{S}$ then coalescence has occurred exactly when $F_{-T,0}(\mathbf{0}) = F_{-T,0}(\mathbf{1})$.

Note that the *Principle* described in Letac (1986, p164) is closely related to the above. However Letac's arguments apply under the condition of almost-sure convergence of the Y_t (rather than the stronger condition of coalescence in finite time) and Letac applies the principle to deliver powerful theoretical characterization arguments rather than simulation algorithms. Letac's ideas have found application in actuarial science as well as theoretical probability: see Dufresne (1991).

I was introduced to Propp and Wilson's idea by Julian Besag at a conference at CWI Amsterdam in November 1995, organized by Ilya Molchanov and Michael Keane. After discussion with Adrian Baddeley, I set myself the challenge of considering how one might make these general ideas work in the context of simulation of point processes, and the following is the first fruits of this investigation. There are two main issues to deal with: firstly that the state-space $\mathcal{S}$ is now uncountable, and secondly that monotonicity does not apply for many interesting point process models. We see below how *both* these issues can be dealt with effectively in the case of perfect simulation for the area-interaction point process. (Indeed the methods described here also generalize to repulsive Strauss point processes, including the limiting case of the excluded-area or hard-core point process.) So this paper records a successful response to this challenge. Clearly stochastic geometers must now be more ambitious about their simulation objectives: after the ideas of Propp and Wilson we can no longer be content merely with long-run approximations to equilibrium distributions but instead must strive after perfection.

It should be mentioned that spatial birth-and-death process equilibrium theory is no stranger to coupling techniques, which figure in the seminal work of Preston (1976), also in Lotwick & Silverman (1981) (for hard-core processes) and Møller (1989) (where a geometric ergodicity result is established). A recent paper by Georgii & Küneth (1996) establishes similar results by direct arguments about point processes. What is new here is the application of coupling to provide exact simulation applicable to several varieties of point process.

It is appropriate in the setting of this volume to note that this work was greatly assisted by the World-Wide Web, which enabled rapid and easy access to a preprint version of Propp & Wilson (1996). Coupling ideas of this kind typically spread extremely rapidly because of their profoundly intuitive nature and (*post hoc*!) simplicity: for example the fundamental reflection Brownian coupling idea due to Lindvall (1982) spread so fast that by the time I heard of it the idea had already become detached from the name of its originator and it took some careful detective work to find out who deserved the credit! Propp and Wilson's idea has the same stamp of profound simplicity; the World-Wide Web has ensured that their names remain attached to it. May modern technology continue to be servant rather than master to scholarship.

Finally I gratefully acknowledge the encouragement and helpful remarks of Professors J. Møller (Aalborg University) and O. Häggström (Chalmers University, Goteborg).

In the following we shall use some conventional notation from stochastic geometry:

m_r denotes Lebesgue measure for the Euclidean space $\mathbb{R}^r$;
$\oplus$ denotes Minkowski addition applied to sets: thus $A \oplus B = \{x : x = a + b \text{ for } a \in A, b \in B\}$;
If X is a finite point process then $\#(X)$ is the number of points in X.

2: The Area-Interaction Point Process X.

This process was introduced in a careful study by Baddeley and Van Lieshout (1995). It is intended to be a single point process model which for differing parameters exhibits either clustering (*attraction*) or dispersion (*repulsion*) between points, relative to the Poisson point process. This is achieved by weighting the Poisson point process distribution according to an exponential of the area (in the planar case) of the union of congruent planar bodies centered at the points of the process. This produces a delightfully simple and appealing model. It should be contrasted with the Strauss point process, commonly used in stochastic geometry, which weights according to an exponential of the number of point pairs closer than some fixed threshold: while the Strauss model is successful in modelling repulsion neither it nor obvious generalizations can produce successful models for attraction (indeed in the case of the Strauss model a divergence means the resulting point process fails to exist!). Thus the importance of the area-interaction point process for stochastic geometry lies exactly in the fact that it provides a simple model which for various parameter values can model either attraction or repulsion between points.

In this section we briefly summarize the main features of this point process, as described by Baddeley and Van Lieshout.

In its simplest form, the area-interaction point process is a random process X of points in $\mathbb{R}^d$, with distribution having a Radon-Nikodým density or weight $p(X)$ with respect to the unit-rate Poisson process restricted to a compact window, where

$$p(X) \quad = \quad \alpha\lambda^{\#(X)}\gamma^{-m_d(X\oplus G)} \, .$$

Here α is a norming constant, λ, γ are positive parameters, and the *grain* G is a compact (typically convex) subset of $\mathbb{R}^d$. The set $X \oplus G$ is given by

$$X \oplus G \quad = \quad \bigcup \{x \oplus G : x \in X\} \, ,$$

so it would be the *Boolean model* formed by grain G and germs from X if X were a standard (*i.e.* unweighted) Poisson process. Thus the realization $X \oplus G$ can be thought of as a "weighted Boolean model".

As is usual in the theory of such point processes, the normalization parameter α is not computable in closed form except in trivial cases: the area-interaction means that it is related to exponential moments of coverage of Boolean models. Of course this intractability means simulation methods are of great importance.

The parameter λ can be thought of as controlling the intensity of X (though in fact the actual intensity of the point process X is a complicated function of λ, γ, and G).

The parameter γ controls the *area-interaction* between the points of X: $\gamma > 1$ is the *attractive* case and $\gamma < 1$ is the *repulsive* case.

The attractive case of $\gamma > 1$ (and grain G a ball) was already known in statistical physics as the "penetrable spheres" model of Widom & Rowlinson (1970), important because (unusually for such models) it is possible to prove that phase-transition occurs in case $d > 1$. Further details of this connection also provide a simulation method for area-interaction point processes *via* accept-reject methods producing single-type margins of conditioned two-type Poisson processes, both for $\gamma > 1$ and $\gamma < 1$.

The area-interaction point process is well-defined for all values of the parameters λ, γ, and compact G, satisfies a stability condition of Ruelle type, and can therefore be extended to a well-defined point process over the whole of Euclidean space.

There are several possible generalizations. One which is relevant here is to replace Lebesgue measure by other Borel measures: replacing Lebesgue measure m_d by a suitable multiple of counting measure on a lattice is convenient for simulation purposes.

As well as methods related to conditioned two-type Poisson point processes as indicated above, one can simulate the area-interaction point process using the common method of simulating spatial birth-and-death processes till they are deemed to be close to equilibrium. In the remainder of this paper we will use this as our test case for perfect simulation of point processes: our objective is to describe a modification of spatial birth-and-death process simulation which does indeed give perfect (and practical!) simulation for an area-interaction point process. We will also see that this modification is susceptible to generalization, for example to Strauss point processes.

3: A coupling construction for perfect simulation of X.

Consider the task of perfect simulation of an attractive area-interaction process X in a bounded window $W \subset \mathbb{R}^d$. Let $G \subset \mathbb{R}^d$ be the compact non-random grain on which the definition of X is based, let λ be the intensity of the underlying reference Poisson process, and let $\gamma > 1$ be the parameter of (attractive) area-interaction.

We shall use a space-time Boolean model and $[0, 1]$-valued marks to build coalescing spatial birth-and-death processes whose common equilibrium distribution is the distribution of X. (These spatial birth-and-death processes provide the analogue of the discrete-time stochastic flow $F_{s,t}$ of Section 1.) The coalescence means that we are able to sample from the spatial birth-and-death processes as if they had been run from time $-\infty$, and consequently our sample must have the equilibrium distribution. A point of finesse is that the coalescence takes place *conditional on the realization of the marked grain process of the Boolean model*: nevertheless the Propp-Wilson argument goes through.

Our construction will assume empty boundary conditions, but non-trivial boundary conditions can be catered for by an easy modification.

The underlying space-time Boolean model is based on grains which are space-time cylinders

$$(G \oplus x) \times [s,\ s + \ell\,] \quad \subset \quad \mathbb{R}^d \times \mathbb{R}$$

with spatial cross-section $G \oplus x$. We say that the *birth time* of the above space-time cylinder is s, while its *death time* is $s + \ell$. We parametrize such space-time cylinders by the representing space built from coordinates x, s, ℓ:

$$(x, s, \ell\,) \quad \in \quad \mathcal{C} \quad = \quad \mathbb{R}^d \times \mathbb{R} \times (0, \infty)\,.$$

We want to mark the grains with independent uniform $[0, 1]$ marks, so let Z be the Poisson process on $\mathcal{C} \times [0, 1]$ governed by the intensity measure

$$\left(\lambda e^{-\ell}\, \mathbb{I}_{[\,x \in W\,]}\right)\ m_d(dx)\, m_1(ds)\, m_1(d\ell)\, m_1(dp)\,.$$

Hence the intensity measure factorizes as the product of the uniform measure of intensity λ on $W \times \mathbb{R}$, the exponential(1) distribution for length ℓ and the uniform $[0, 1]$ distribution for mark p. From Z we derive a space-time Boolean model $\Psi \subset W \times \mathbb{R}$ by taking the union of the grains and forgetting the marks:

$$\Psi \quad = \quad \bigcup \left\{ (G \oplus x) \times [s, s\ + \ell\,] \ :\ [\,(x, s, \ell\,);\ p\,] \in Z \right\}.$$

Now fix $-T < 0$: for $u \in [-T,\ 0]$ we define time-evolving point processes $Y^{\max}(-T, u)$, $Y^{\min}(-T, u)$, $Y(-T, u)$ on $W \subseteq \mathbb{R}^d$. They are grown recursively from initial configurations $Y^{\max}(-T, -T)$, $Y^{\min}(-T, -T)$, $Y(-T, -T)$, and will be built using the sub-process of marked points representing those space-time cylinders of Z which overlap $[-T,\ 0]$:

$$Z^{[-T,\ 0]} \quad = \quad \left\{ [\,(x, s, \ell\,);\ p\,] \in Z \ :\ -T \leq s + \ell,\ s \leq 0 \right\}.$$

The initial configurations are defined as follows: set

$$\begin{aligned}
Y^{\max}(-T, -T) \quad &= \quad \left\{ x \ :\ [\,(x, s, \ell\,);\ p\,] \in Z^{[-T,\ 0]},\ s \leq -T \leq s + \ell \right\} \\
Y^{\min}(-T, -T) \quad &= \quad \left\{ x \ :\ [\,(x, s, \ell\,);\ p\,] \in Z^{[-T,\ 0]},\ s \leq -T \leq s + \ell,\ p \leq \gamma^{-m_d(G)} \right\}
\end{aligned}$$

while $Y(-T, -T)$ is some arbitrary subset of $Y^{\max}(-T, -T)$ which we do not specify, except to require that it must be a superset of $Y^{\min}(-T, -T)$. We shall call such a $Y(-T, u)$ ($u \in [-T, 0]$) *appropriate*, since it is not necessary to consider any other Y in order to obtain the desired equilibrium.

It will be convenient to define three auxiliary spatial germ-grain models $\Xi^{\max}(-T, u)$, $\Xi^{\min}(-T, u)$, $\Xi(-T, u)$ by

$$\begin{aligned}
\Xi^{\max}(-T, u) \quad &= \quad Y^{\max}(-T, u) \oplus G \quad = \quad \bigcup \left\{ G \oplus x \ :\ x \in Y^{\max}(-T, u) \right\}, \\
\Xi^{\min}(-T, u) \quad &= \quad Y^{\min}(-T, u) \oplus G \quad = \quad \bigcup \left\{ G \oplus x \ :\ x \in Y^{\min}(-T, u) \right\}, \\
\Xi(-T, u) \quad &= \quad Y(-T, u) \oplus G \quad = \quad \bigcup \left\{ G \oplus x \ :\ x \in Y(-T, u) \right\}.
\end{aligned}$$

Now let $-T < t_1 < \ldots < t_N < 0$ be the ordered list of space-time cylinder birth and death times lying in $[-T,\ 0]$, so that

$$\{t_1, \ldots, t_N\} \quad = \quad [-T,\ 0] \cap \bigcup \left\{ s, s + \ell \ :\ [\,(x, s, \ell\,);\ p\,] \in Z^{[-T,\ 0]} \right\}.$$

By simple properties of Poisson processes we may argue that: no two birth/death times coincide; moreover corresponding to each birth/death time t_i there is a unique $[\,(x, s, \ell\,);\ p\,] \in Z$ with $s = t_i$ in case of birth, and $s + \ell = t_i$ in case of death.

We can now give the recursive definition of $Y(-T, u)$ for $u \in [-T,\ 0]$; the definitions of $Y^{\max}(-T, u)$, $Y^{\min}(-T, u)$ are entirely analogous, but in (c) below replace Ξ by $\Xi^{\max}$, $\Xi^{\min}$ respectively.

Proposition 1:

The following algorithm simulates a spatial birth-and-death process:

(a) *Suppose that $Y(-T,u)$ is already defined, and that t_i is the least birth/death time greater than u. We set*

$$Y(-T,u+v) = Y(-T,u) \qquad \text{for } u \leq u+v < t_i\,.$$

(If there is no such time then we set $Y(-T,u+v) = Y(-T,u)$ for all $v \geq 0$, and the definition of $Y(-T,\cdot)$ is then complete!) What happens next depends on whether t_i is a death time (in which case we go to (b) below) or a birth time (in which case we go to (c) below).

(b) *If t_i is a* **death time**, *so $t_i = s + \ell$ for some $[\,(x,s,\ell\,);\ p\,] \in Z$, then we* **delete** *the point x: we set*

$$Y(-T,t_i) = Y(-T,t_i-) \setminus \{x\}\,.$$

(Here, of course, $Y(-T,t_i-)$ is the configuration of $Y(-T,\cdot)$ immediately prior to time t_i.) We then continue the construction using (a) above.

(c) *If t_i is a* **birth time**, *so $t_i = s$ for some $[\,(x,s,\ell\,);\ p\,] \in Z$, then we test whether*

$$p \quad \leq \quad \gamma^{-m_d((G\oplus x)\setminus\Xi(-T,t_i-))}\,.$$

(Recall that $\Xi(-T,t_i-) = \Xi(-T,t_i-) \oplus G = \bigcup\{G\oplus x\ :\ x \in Y(-T,t_i-)\}$.) If this is the case then we **add** *the point x: we set*

$$Y(-T,t_i) = Y(-T,t_i-) \cup \{x\}\,.$$

Otherwise we **do nothing**: *we set*

$$Y(-T,t_i) = Y(-T,t_i-)\,.$$

In either case we then continue the construction using (a) above.

It is clear that this construction is complete and unambiguous (because to each birth/death time there corresponds a unique representing point of Z as mentioned above) and can be programmed on a computer. Indeed I have written a `C` program `Perfect`, currently in the experimental stage, which uses the above algorithm to conduct perfect simulations of area-interaction planar point processes, subject to an approximation of Lebesgue measure by a discrete measure as mentioned in Section 2 above. Further information about `Perfect` can be obtained from me at `w.s.kendall@warwick.ac.uk`.

In order to see how this provides perfect simulation it is necessary only to check a number of simple points, which we list below as propositions and whose proofs are largely left as exercises for the reader. (As is often the case with coupling arguments, to give heavily rigorous proofs would be rather like explaining a joke: it tends to spoil the fun! However the following propositions really deserve arguments demanding more space than can be spared here: I hope to deliver these arguments in a follow-up article dealing with further work.) As a consequence of these propositions one can obtain perfect simulations of area-interaction point processes from inspection of simulation runs of these spatial birth-and-death processes.

Proposition 2:

The spatial birth-and-death processes $Y^{\max}(-T,u)$, $Y^{\min}(-T,u)$, $Y(-T,u)$, viewed as functions of $u \in [-T,\ 0]$, have the distribution of the target area-interaction process X as equilibrium distribution.

To check this is merely a matter of confirming that detailed balance holds and that simple regularity conditions apply, since in equilibrium these spatial birth-and-death processes are reversible. See Preston (1977) for the theory of spatial birth-and-death processes and Ripley (1977), Ripley (1979) for the use of spatial birth-and-death processes in simulation of point processes such as X. Stoyan, Kendall & Mecke (1995, Chapter 5) provide a summary exposition in book form. Section 3 of Baddeley and Van Lieshout (1995) gives details which apply to the processes $Y^{\max}(-T,\cdot)$, $Y^{\min}(-T,\cdot)$, $Y(-T,\cdot)$.

Proposition 3:

The following simple modification of the algorithm in Proposition 1 deals with the repulsive case of $\gamma < 1$: increase the intensity of Z to

$$\left(\gamma^{-m_d(G)} \times \lambda e^{-\ell}\,\mathbb{I}_{[\,x \in W\,]}\right)\; m_d(dx)\, m_1(ds)\, m_1(d\ell)\, m_1(dp)\,,$$

alter the initialization of $Y^{\min}(-T,T)$ to

$$Y^{\min}(-T,-T) \quad = \quad \left\{ x \;:\; [\,(x,s,\ell\,);\; p\,] \in Z^{[-T,\;0]},\; s \leq -T \leq s+\ell,\, p \leq \gamma^{m_d(G)} \right\}\,,$$

and alter the test in (c) to **add** *x when*

$$p \quad \leq \quad \gamma^{m_d(G) - m_d((G \oplus x) \setminus \Xi(-T, t_i -))}\,.$$

The result of the construction is then a spatial birth-and-death process whose equilibrium distribution is that of a repulsive area-interaction point process of parameter $\gamma < 1$ based on a Poisson point process of underlying intensity λ. (Again this is simply a matter of confirming that detailed balance holds and that simple regularity conditions apply.) As we see below, it is then less clear how one should confirm that exact simulation has taken place, since this is no longer an attractive point process. Nevertheless we will show that it can be done.

The construction of the processes $Y^{\max}(-T,u)$, $Y^{\min}(-T,u)$, $Y(-T,u)$ is based entirely on the randomness provided by $Z^{[-T,\;0]}$. For $-S < -T$ the point process $Z^{[-S,\;0]}$ can be produced by augmenting $Z^{[-T,\;0]}$, and the resulting construction of $Y^{\max}(-S,u)$, $Y^{\min}(-S,u)$, $Y(-S,u)$ is strongly related to the construction over $[-T,\;0]$: for example for all S the final patterns $Y^{\max}(-S,0)$, $Y^{\min}(-S,0)$, $Y(-S,0)$ are subsets of a single fixed point pattern

$$\left\{ x \;:\; [\,(x,s,\ell\,);\; p\,] \in Z,\; s \leq 0 \leq s+\ell \right\}.$$

In particular we can consider *all* constructions based on a single given realization of the underlying space-time cylinder process Z. In the attractive case ($\gamma > 1$) this leads to a crucial monotonicity property.

Proposition 4:

In the attractive case $\gamma > 1$, if $-S \leq -T \leq u \leq 0$ then

$$Y^{\min}(-T,u) \;\subseteq\; Y^{\min}(-S,u) \;\subseteq\; Y(-S,u) \;\subseteq\; Y^{\max}(-S,u) \;\subseteq\; Y^{\max}(-T,u)$$

whenever $Y^{\min}(-S,-S) \subseteq Y(-S,-S) \subseteq Y^{\max}(-S,-S)$.

This is best seen by arguing recursively. Certainly the result holds for $u = -T$, since the algorithm described in Proposition 1 always adds points $[(x,s,\ell\,);\; p\,] \in Z$ for which $p \leq \gamma^{-m_d(G)}$. If the result holds up to time u then we have a similar inclusion for the corresponding random sets:

$$\Xi^{\min}(-T,u) \;\subseteq\; \Xi^{\min}(-S,u) \;\subseteq\; \Xi(-S,u) \;\subseteq\; \Xi^{\max}(-S,u) \;\subseteq\; \Xi^{\max}(-T,u)\,.$$

But if $[(x,u,\ell\,);\; p\,] \in Z$ is a point to be added at time u then this monotonicity, together with the birth criterion

$$p \quad \leq \quad \gamma^{-m_d((G \oplus x) \setminus \Xi^*(-T, t_i -))}\,,$$

for Ξ^* running through $\Xi^{\min}(-T,u)$, $\Xi^{\min}(-S,u)$, $\Xi(-S,u)$, $\Xi^{\max}(-S,u)$, $\Xi^{\max}(-T,u)$, implies that x will be added in such a way as to maintain monotonicity. (Remember that we are dealing with the attractive case $\gamma > 1$ here.)

Thus appropriate $Y(-S,\cdot)$ form a subclass of appropriate $Y(-T,\cdot)$. Unfortunately in the repulsive case $\gamma < 1$ the above argument breaks down badly and this monotonicity does not hold. We shall see below (Proposition 7) how to fix this in a practical way. First we note how it is that coalescence allows us to do perfect simulation.

Suppose that we are given an initial point pattern $\{x_1, ..., x_k\}$ with which to start off the simulation at time $-S$. Point x_i lasts for an exponential time r_i (of unit mean), so all initial points die out by time $-T$ if $-S < -T - \max_i r_i$. Using the space-time cylinder process Z, we construct a spatial birth-and-death process $\tilde{Y}(-S, u)$ with initial condition $\{x_1, ..., x_k\}$ at time $-S$ as follows. We set $\tilde{\Xi}(-S, u) = \tilde{Y}(-S, u) \oplus G$ for $-S \leq u \leq 0$, and set $\tilde{Y}(-S, -S) = \{x_1, ..., x_k\}$. The point x_i is deleted from the pattern at time $-S + r_i$. Otherwise $\tilde{Y}(-S, \cdot)$ evolves according to the algorithm described in Proposition 1, using of course $\tilde{\Xi}(-S, \cdot)$ to compute the area bias and so to determine when points from Z are to be added to the pattern.

It follows from the algorithm that for $-T \in [-S + \max_i r_i, 0]$ the point pattern $\tilde{Y}(-S, -T)$ is *appropriate* in the sense defined above: that is to say $Y^{\min}(-T, -T) \subseteq \tilde{Y}(-S, -T) \subseteq Y^{\max}(-T, -T)$. Consequently if T is large enough for coalescence, so that $Y^{\min}(-T, 0) = Y^{\max}(-T, 0)$, then $Y^{\min}(-T, 0) = Y^{\max}(-T, 0) = \tilde{Y}(-S, 0)$ no longer depends on S (as long as $S > T + \max_i r_i$). Hence

(a) $\lim_{S\to\infty} \tilde{Y}(-S, 0)$ exists and has the equilibrium distribution,

(b) by definition of coalescence being observable, we can choose a "stopping time" T (that is to say, $-T_0 \leq -T$ is observable from the realization $Z^{[-T_0, 0]}$) such that $Y(-T, 0) = \lim_{S\to\infty} \tilde{Y}(-S, 0)$ for all appropriate $Y(-T, 0)$, including $Y^{\min}(-T, 0)$.

Thus *in principle* we can sample from $\lim_{S\to\infty} \tilde{Y}(-S, 0)$ and thus from the equilibrium distribution, which is the desired distribution of the area-interaction point process. A simple percolation argument for the Boolean model Ψ in the bounded window W now shows that for sufficiently large T we have $Y^{\min}(-T, 0) = Y^{\max}(-T, 0)$, and therefore we have coalescence and the possibility at least in principle of perfect simulation.

For suppose that $-\tau$ is the random time which is the supremum of times such that the fixed-time slice $\mathbb{R}^d \times \{0\}$ is *not* topologically connected to $\mathbb{R}^d \times \{-T\}$ by Ψ whenever $-T < -\tau$. (So $\mathbb{R}^d \times \{0\}$ and $\mathbb{R}^d \times \{0\}$ are in different connected components of $\mathbb{R}^d \times \{0\} \cup \Psi \cup \mathbb{R}^d \times \{1\} \subset \mathbb{R}^d \times \mathbb{R}$ whenever $-T < -\tau$.) This is almost surely finite since we have stipulated that the window W must be bounded. Notice also that τ is a "stopping time", in the sense that the event $[-T < -\tau]$ is measurable with respect to $Z^{[-T,\, 0]}$, so that τ is observable based only on simulation over the range $[-T, 0]$. It follows easily from the construction that if $-T < -\tau$ then the final values $Y^{\max}(-T, -T)$, $Y^{\min}(-T, -T)$, $Y(-T, -T)$ must all agree and no longer depend on T or initial conditions. Consequently the construction can be used to make a perfect simulation of an area-interaction process:

Proposition 5:

Perfect simulation of an area-interaction point process may be obtained by the following **algorithm***: Choose a $-T_1 < 0$, simulate and construct $Z^{[-T_1,\, 0]}$, $Y^{\max}(-T_1, 0)$. If $-T_1 < -\tau$ then $Y^{\max}(-T_1, 0)$ ($= Y(-T_1, 0)$ for all appropriate $Y(-T_1, 0)$) is a sample from the desired equilibrium distribution. If not then choose $-T_2 < -T_1 < 0$, etc, and repeat till $-T_n < -\tau$. At each stage care must be taken to extend the fundamental space-time cylinder process Z so that $Z^{[-T_{n-1}, 0]} \subseteq Z^{[-T_n, 0]}$. Then return $Y^{\max}(-T_n, 0)$ as a sample from the desired equilibrium distribution.*

The required extension property $Z^{[-T_{n-1}, 0]} \subseteq Z^{[-T_n, 0]}$ was easily obtained in the `C` implementation `Perfect` by using pointers to build up a list of the space-time cylinders in $Z^{[-T_{n-1}, 0]}$, respectively $Z^{[-T_n, 0]}$, deriving a list of potential birth/death times, then sorting the derived list in time order in order to conduct the actual simulation.

This provides perfect simulation for area-interaction point processes whether attractive or repulsive. However τ may be large, and so we need more practical methods of determining when coalescence might have happened. Proposition 4 provides the key for the attractive case $\gamma > 1$, as monotonicity gives an easy way of avoiding checking for $-T_n < -\tau$:

Proposition 6:

Perfect simulation of an area-interaction point process may be obtained by the following **modified algorithm***: Increase T in the construction given in Proposition 5 (always by extending previous constructions of the fundamental space-time cylinder process Z) until it is observed that*

$$Y^{\max}(-T_n, 0) \quad = \quad Y^{\min}(-T_n, 0)\,.$$

Then $Y^{\max}(-T_n, 0) = Y^{\min}(-T_n, 0)$ is a sample from the desired equilibrium distribution.

For in such a case Proposition 4 implies that $Y(-S,0) = Y^{\min}(-T,0) = Y^{\max}(-T,0)$ for all appropriate $Y(-S,0)$ where $-S < -T$, even if $-T_n < -\tau$ does not yet hold, and so coalescence has occurred. Simulation experiments have shown that this monotonicity idea (already used to great effect in the case of finite discrete state-space by Propp & Wilson, 1996) provides a practical means of perfect simulation of an attractive area-interaction spatial birth-and-death process (in contrast the algorithm described in Proposition 5 is unusably slow): for example the pattern of 55 points in Figure 1, with parameters as described there, can be produced in 30 seconds by the (unoptimized) experimental version of `Perfect`, working in a slow multi-tasking environment using less than 640k RAM on my five-year-old personal computer (Acorn A440/1 `ARM3` with `FPA10` floating-point coprocessor). This simulation used 495 birth-or-death steps to produce the 55 points of $Y^{\min}(-T,0) = Y^{\max}(-T,0)$ (incidentally it is interesting to compare this with the *heuristic* recommendations which can be found in the literature, suggesting for example $4\times$ number of points for moderate parameter values for systematic single-point updating algorithms for Strauss point processes).

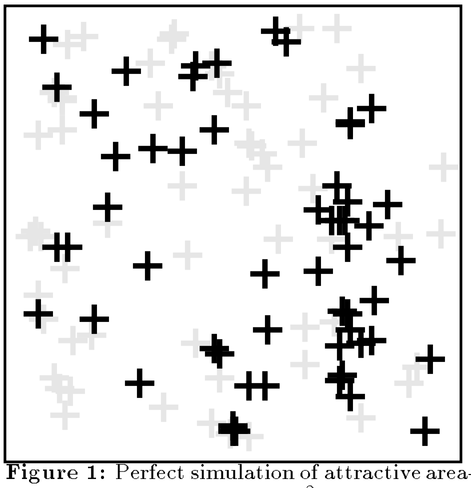

Figure 1: Perfect simulation of attractive area-interaction point process: 55 points, $\gamma = e^2 = 7.39$, $\lambda = 1$, square window of side-length 10, square grain side-length 0.7. Shaded crosses mark centres of sections of space-time cylinders in Z which intersect the time-slice 0.

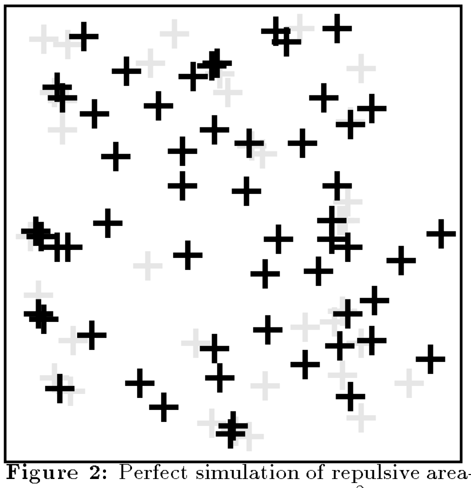

Figure 2: Perfect simulation of repulsive area-interaction point process: 55 points, $\gamma = e^{-2} = 0.135$, $\lambda = 0.25$, square window of side-length 10, square grain side-length 0.7. Shaded crosses mark centres of sections of space-time cylinders in the modified Z which intersect the time-slice 0.

Monotonicity is not available for the repulsive case, and simulation experiments show that the percolation criterion using τ defined in Proposition 5 is not practical. Indeed my joint PhD student Elke Thönnes has observed that the probability $\mathbb{P}\left[-\tau > -T\right]$ is subject to a crude bound in terms of the distribution of residual busy period for an $M/M/\infty$ queue, but that this is unusably large in practical situations. Better bounds will be obtained from continuum percolation arguments: in particular even when the observation window is unbounded we still know that τ will be almost surely finite whenever the underlying intensity λ (respectively $\tilde{\lambda}\gamma^{-m_d(G)}$ for the repulsive case) is less than a percolation constant. Indeed rather similar percolation arguments have been used by Chayes, Chayes & Kotecký (1995) to give elegantly simple derivations of phase transition for the attractive area-interaction point process, in its older guise as the Widom-Rowlinson model. From the point of view of perfect simulation, however, theoretical bounds on $\mathbb{P}\left[-\tau < -T\right]$ *etc* are of lesser importance than the *observability* of $[-\tau < -T_n]$, which suggests that perfect simulation might be a feasible possibility.

It is delightful to report that there is a remarkably simple modification of the algorithm described in Propositions 1, 5, and 6 which is effective in determining equilibrium in the repulsive case.

Proposition 7:

Modify the constructions of $Y^{\max}$, $Y^{\min}$ from the method described in Proposition 1 as follows: in (c) of the construction, when t_i is a birth time, the corresponding point x is added in the case of $Y^{\max}$ when

$$p \quad \leq \quad \gamma^{m_d(G) - m_d\left((G\oplus x)\setminus \Xi^{\min}(-T,t_i-)\right)},$$

and in the case of $Y^{\min}$ when

$$p \quad \leq \quad \gamma^{m_d(G)-m_d((G\oplus x)\setminus\Xi^{\max}(-T,t_i-))}\,.$$

Then the algorithm which samples $Y^{\max}(-T,0) = Y^{\min}(-T,0)$ after coalescence produces a sample from the relevant repulsive area-interaction point process.

In this algorithm the constructions of $Y^{\max}$, $Y^{\min}$ are now coupled, and are no longer individually Markov. (In order to keep the Markov property we would have to consider the two-type spatial birth-and-death process $(Y^{\max}, Y^{\min})$.) But it is now simple to check by an inductive argument that

$$Y^{\min}(-T,u) \quad \subseteq \quad Y(-S,u) \quad \subseteq \quad Y^{\max}(-T,u)$$

whenever $-S \leq -T$, $u \in [-T,0]$, and $Y^{\min}(-T,-T) \subseteq Y(-T,-T) \subseteq Y^{\max}(-T,-T)$. For the desired relation holds at time $u=-T$, deaths are the same, and the test for a $Y^{\max}$ birth (being based on $Y^{\min}$) is always no more stringent than the test for the corresponding $Y^{\min}$ birth (being based on $Y^{\max}$ which by induction contains $Y^{\min}$), nor indeed is it more stringent than the test for the corresponding Y birth. It therefore follows that if $Y^{\min}(-T,0) = Y^{\max}(-T,0)$ then $Y(-S,0) = Y^{\min}(-T,0) = Y^{\max}(-T,0)$ for all appropriate $Y(-S,0)$ where $-S < -T$, and so the common pattern of $Y^{\min}(-T,0) = Y^{\max}(-T,0)$ is a sample from the equilibrium distribution. In summary, we replace the lower and upper bounding spatial birth-and-death processes of the attractive case with well-behaved lower and upper bounding processes which are no longer spatial birth-and-death processes, but which are well-adapted to our problem!

Simulation experiments have shown that this provides a practical means of perfect simulation of a repulsive area-interaction spatial birth-and-death process: for example the pattern of 55 points in Figure 2, with parameters as described there, can be produced in 7 seconds by `Perfect`, running on the same configuration as above. (That the same number of points were obtained as in the Figure 1 is the result of an informal rejection-sampling undertaken in order to produce comparable figures!) This simulation used 1132 birth-or-death steps to produce the 55 points of $Y^{\min}(-T,0) = Y^{\max}(-T,0)$.

Implementation of the repulsive case required the addition of just 3 statements in `C` code!

The case of non-empty boundary conditions can be dealt with quite simply for all variants of the algorithm described above, just by adding the boundary points to both $Y^{\min}$ and $Y^{\max}$ for all time.

4: Conclusion.

The above shows that perfect simulation is possible for both attractive and repulsive area-interaction point processes. Of course this is just the beginning! Further questions which should now be addressed include

(a) Can one establish good bounds on the run-times required by these perfect convergence algorithms? A combination of heuristics and crude simulation experiments suggests that if L is the linear dimension of (square) window size, then *for fixed parameters which are not too extreme* the time required for simulations over windows of increasing size should be about $O\left(L^2(\log(L))^2\right)$ (the dominant contribution here arising from sorting birth/death incidents arising from the underlying space-time cylinder process over the time interval required for coupling). Whatever the correct asymptotic, the constant multiplier will be extremely sensitive to the parameters. In particular it is interesting to note that the repulsive case $\gamma < 1$ actually results in faster convergence than the attractive case $\gamma > 1$, if one standardizes on observed intensity of points. It would be useful to have a mathematical explanation of this.

(b) The monotonicity of this construction can be exploited to prove monotonic relationships between various area-interaction point processes, and also limit theorems. The simplest example of this is that every area-interaction point process can be placed on a probability space such that it is contained in one Poisson point process, and contains another! Notice also that for fixed λ and for $\gamma > 1$ the construction gives a simple proof that the intensity of the area-interaction point process decreases as γ increases. Finally, for fixed γ the construction shows that for high λ (hence high intensity) the area-interaction point process is nearly Poisson. None of these results is particularly deep (indeed this monotonicity has already been exploited several times in the literature: see Chayes, Chayes & Kotecký (1995) for a recent example, and Georgii & Küneth (1996) for non-coupling proofs), but they illustrate the point that a good simulation method provides qualitative as well as quantitative insight into a process!

(c) As with the standard spatial birth-and-death simulation approach to point processes, there is an alternative algorithm for perfect simulation based on constant birth rate and variable death rate.

(d) Is there a way of doing perfect simulation for add-delete algorithms for a wide variety of point processes?

(e) Can one extend these ideas to other point processes, and more generally to the situation of Markov chain Monte Carlo sampling (MCMC) from posteriors obtained by conditioning point processes on observed discrete data? Certainly there are simple modifications of the algorithm described in Proposition 7 (which can be implemented by only minor changes in the `C` program `Perfect`) yielding perfect simulation of Strauss point processes with purely repulsive pairwise interactions (in other words, the amenable case typically arising in spatial statistics contexts), and even excluded volume point processes (but here theoretical considerations, which I have confirmed by simulation experiments, show that the method slows down rapidly as the underlying Poisson intensity approaches a certain directed continuum percolation threshold). More generally, any Markov point process with bounded local energy should be amenable to this approach: this includes (for appropriate parameter values) the connected-component interaction ("continuum random cluster model") proposed by Møller. I plan to follow up on this in a further paper. The question of what can be learned for MCMC seems much harder to me: maybe all one can obtain will be useful diagnostics based on the coupling ideas described above.

(f) A particularly pleasing aspect of such extensions is that purely repulsive Strauss point processes and excluded-area point processes can be simulated using the same underlying space-time cylinder process Z as is used for the area-interaction point processes. This allows for both theoretical and experimental comparisons between rather different kinds of point processes.

(g) It would be useful and possibly of theoretical interest to work out faster ways of determining when convergence has occurred for excluded-area point processes. When viewed in space-time the simulation algorithm above can be regarded as a feed-forward logical neural network, with threshold interaction and combining inputs using inclusive `OR`. Unfortunately this approach looks infeasible, because of links between coalescence-determination and the NP-complete `SATISFIABILITY` problem of computer science (Garey & Johnson, 1979, page 39). However perhaps the intrinsic localization of the spatial birth-and-death process can come to the rescue here.

Such questions place a premium on examining and inventing further varieties of coupling methods. For example, consider the challenge, *could reflection coupling be useful in some cases where monotonicity does not apply?* It might be possible to modify the geometric ideas of Kendall (1986) to produce practical algorithms which in suitable cases recognize (at least approximately) when the final result does not depend on the initial conditions. Experience in stochastic differential geometry suggests that some successful coupling strategies work in rather counter-intuitive situations; see Ben Arous, Cranston & Kendall (1995). There is some work on nonhomeomorphic coalescing stochastic flows (Darling, 1987) which may be relevant here. In any case, the success of the preliminary work reported above in providing perfect simulation of area-interaction point processes, and particularly the surprisingly minor modification which deals effectively with repulsive area-interaction point processes, encourages me in the view that creativity rather than pessimism is the appropriate response to this challenge.

References

Baddeley, A.J. & M.N.M. van Lieshout (1995), "Area-interaction point processes," *Ann. Inst. Statist. Math.* **47**, 601–619.

Ben Arous, G., M. Cranston & W.S. Kendall (1995), "Coupling constructions for hypoelliptic diffusions: Two examples," *Proceedings of Symposia in Pure Mathematics* **57**, 193–212.

Besag, J. & P.J. Green (1993), "Spatial statistics and Bayesian computation (with discussion)," *J. Roy. Statist. Soc. Ser. B* **55**, 25–37.

Chayes, J.T., L. Chayes & R. Kotecký (1995), "The analysis of the Widom-Rowlinson model by stochastic geometric methods," *Comm. Math. Phys.* **172**, 551–569.

Chen, M.-F. (1994), "Optimal couplings and application to Riemannian geometry," in *Probability theory and Mathematical Statistics*, B. Grigelionis *et al.*, ed. **1**, VPS / TEV, 15 pp.

Darling, R.W.R. (1987), "Constructing nonhomeomorphic stochastic flows," *Mem. Amer. Math. Soc.* **376**.

Dufresne, D. (1991), "The distribution of a perpetuity, with applications to risk theory and pension funding," *Scand. Actuarial J.* 39-79.

Garey, M.R. & D.S. Johnson (1979), in *Computers and Intractability: A guide to the theory of NP-completeness*, W. H. Freeman and Company, San Francisco, CA.

Georgii, H.-O. & T. Küneth (1996), "Stochastic comparison of point random fields," (submitted).

Kendall, W.S. (1986), "Nonnegative Ricci curvature and the Brownian coupling property," *Stochastics* **19**, 111–129.

Letac, G. (1986), "A contraction principle for certain Markov chains and its applications," *Contemporary mathematics* **50**, 263–273.

Lindvall, T. (1982), *On Coupling of Brownian Motions.* Technical report 1982-23, Department of Mathematics, Chalmers University of Technology & University of Göteborg, Göteborg.

Lotwick, H.W. & B.W. Silverman (1981), "Convergence of spatial birth-and-death processes," *Math. Proc. Cambridge Philos. Soc.* **90**, 155–165.

Møller, J. (1989), "On the rate of convergence of spatial birth-and-death processes," *Ann. Inst. Statist. Math.* **41**, 565–581.

Preston, C.J. (1976), "Random fields," *Lect. Notes in Math.* **534**.

Preston, C.J. (1977), "Spatial birth-and-death processes," *Bull. Inst. Internat. Statist.* **46 (2)**, 371–391.

Propp, J.G. & D.B. Wilson (1996), "Exact sampling with coupled Markov chains and applications to statistical mechanics," *Random Structures and Algorithms* To appear.

Ripley, B.D. (1977), "Modelling spatial patterns (with discussion)," *J. Roy. Statist. Soc. Ser. B* **39**, 172–212.

Ripley, B.D. (1979), "Simulating spatial patterns: dependent samples from a multivariate density," *Applied Statistics* **28**, 109–112.

Stoyan, D., W.S. Kendall & J. Mecke (1987, Second edition 1995), *Stochastic geometry and its applications*, Wiley, Chichester (First edition joint with Akademie Verlag, Berlin).

Widom, B. & J.S. Rowlinson (1970), "New model for the study of liquid-vapor phase transitions," *J. Chem. Phys.* **52**, 1670–1684.

W.S. Kendall, Department of Statistics, University of Warwick, Coventry CV4 7AL, UK.
Email: `w.s.kendall@warwick.ac.uk`

Other University of Warwick Department of Statistics Research Reports
authored or co–authored by W.S. Kendall.

161: The Euclidean diffusion of shape.
162: Probability, convexity, and harmonic maps with small image I: Uniqueness and fine existence.
172: A spatial Markov property for nearest–neighbour Markov point processes.
181: Convexity and the hemisphere.
202: A remark on the proof of Itô's formula for C^2 functions of continuous semimartingales.
203: Computer algebra and stochastic calculus.
212: Convex geometry and nonconfluent Γ-martingales I: Tightness and strict convexity.
213: The Propeller: a counterexample to a conjectured criterion for the existence of certain convex functions.
214: Convex Geometry and nonconfluent Γ-martingales II: Well–posedness and Γ-martingale convergence.
216: *(with E. Hsu)* Limiting angle of Brownian motion in certain two–dimensional Cartan–Hadamard manifolds.
217: Symbolic Itô calculus: an introduction.
218: *(with H. Huang)* Correction note to "Martingales on manifolds and harmonic maps."
222: *(with O.E. Barndorff-Nielsen and P.E. Jupp)* Stochastic calculus, statistical asymptotics, Taylor strings and phyla.
223: Symbolic Itô calculus: an overview.
231: The radial part of a Γ-martingale and a non-implosion theorem.
236: Computer algebra in probability and statistics.
237: Computer algebra and yoke geometry I: When is an expression a tensor?
238: *Itovsn3*: doing stochastic calculus with *Mathematica*.
239: On the empty cells of Poisson histograms.
244: *(with M. Cranston and P. March)* The radial part of Brownian motion II: Its life and times on the cut locus.
247: Brownian motion and computer algebra (Text of talk to BAAS Science Festival '92, Southampton Wednesday 26 August 1992, with screenshots of illustrative animations).
257: Brownian motion and partial differential equations: from the heat equation to harmonic maps (Special invited lecture, 49th session of the ISI, Firenze).
260: Probability, convexity, and harmonic maps II: Smoothness via probabilistic gradient inequalities.
261: *(with G. Ben Arous and M. Cranston)* Coupling constructions for hypoelliptic diffusions: Two examples.
280: *(with M. Cranston and Yu. Kifer)* Gromov's hyperbolicity and Picard's little theorem for harmonic maps.
292: Perfect Simulation for the Area-Interaction Point Process.
295: On some weighted Boolean models.
296: A diffusion model for Bookstein triangle shape.

If you want copies of any of these reports then please mail your requests to the secretary of the Department of Statistics, University of Warwick, Coventry CV4 7AL, UK, or use email (`L.K.Butterwick@warwick.ac.uk`).