Perfect simulation for correlated Poisson random variables conditioned to be positive

by

Y. Cai and W.S. Kendall

June 1999

Submitted draft: 1.19.
The source is /home/bach/wsk/ms/Cai/perfect.tex (unix)
It was last edited on 14:45:00, 28/06/1999.

Perfect simulation for correlated Poisson random variables conditioned to be positive

Yuzhi Cai and Wilfrid S. Kendall

Department of Statistics, University of Warwick
Coventry CV4 7AL, UK

July 27, 1999

Abstract

In this paper we present a perfect simulation method for obtaining perfect samples from collections of correlated Poisson random variables conditioned to be positive. We show how to use this method to produce a perfect sample from a Boolean model conditioned to cover a set of points: in [11] this special case was treated in a more complicated way. The method is applied to several simple examples where exact calculations can be made, so as to check correctness of the program using χ^2-tests, and some small-scale experiments are carried out to explore the behaviour of the conditioned Boolean model.

1 Introduction

Markov chain Monte Carlo (McMC) provides the opportunity to carry out simulation-based statistical inference, by allowing one to sample *approximately* from a target distribution realised as the equilibrium of a Markov chain. The approximate nature of McMC arises because the output is subject to initialization bias. Recent work on perfect simulation [4, 13] shows that in favourable circumstances McMC simulation algorithms can be modified so as to replace the approximation by an exact simulation procedure.

In particular the *Coupling from the Past* (CFTP) method due to Propp & Wilson [13] has been developed intensively since its discovery in 1994. For example, *dominated CFTP* is a variant of CFTP developed in [8], which permits perfect simulation of a general class of point processes (see also [10]). Dominated CFTP has an importance going beyond point processes, since it allows application of CFTP to a wide class of non-uniformly ergodic Markov chains which need not be monotonic. Kendall & Thönnes [11] used dominated CFTP to carry out perfect simulation in stochastic geometry, showing how to get perfect samples from Boolean models conditioned to cover a finite set of points.

The purpose of this paper is to explain how to get a perfect sample from collections of correlated Poisson random variables conditioned to be positive (briefly, correlated Poisson random variables arise as sums of a basic set of independent Poisson random variables). Following Kendall & Thönnes [11], the motivating application is to obtain a perfect sample from a Boolean model conditioned to cover a set of points. However the construction in [11] is very involved. Conditional Boolean model can be viewed as arising from correlated conditioned Poisson random variables, and this leads to a simpler simulation algorithm; the more general random variables set-up draws attention away from the geometry (which here is a distracting complexity). As well as being of interest to stochastic geometers wishing to investigate the conditional Boolean model, the results reported on here add to a sequence of case-studies [6, 7, 8, 11, 10] aimed at developing CFTP by using it in challenge-problems from stochastic geometry. The particular challenge of perfect simulation of conditional Boolean models has resulted in the formulation of a rather general context for CFTP described as "extended state-space CFTP" below.

The construction of the paper is as follows. In §2 we introduce the notion of correlated Poisson random variables, describe a McMC method for obtaining an approximate sample from the conditional correlated Poisson distribution, and then specify an algorithmic scheme for obtaining a perfect sample, in the context of extended state-space CFTP. In common with most other practical CFTP algorithms this is based on the construction of interleaved sequences of upper- and lower-sandwich processes, bounding intermediate Markov processes which actually have the required equilibrium distribution. Fundamental properties of our CFTP algorithm are established by Theorem 1 in §3: namely that the algorithm will terminate in finite time and that the output has the desired distribution. In §4 we describe the relationship between conditional Boolean models and collections of correlated Poisson random variables conditioned to be positive and thus show how to use our algorithm to conduct CFTP for the conditional Boolean model. Several examples are described in §5. For these examples, we derive explicit conditional distribution functions of the correlated Poisson random variables, enabling us to carry out χ^2-tests in order to check the correctness of the CFTP algorithm. Relationships between the intensity of the germ process, radius of the typical grain, number of conditioning points and the coalescence time are investigated in §6. Finally, some comments and conclusions are presented in §7.

2 Correlated Poisson random variables conditioned to be positive

We begin this section by introducing the concept of collections of correlated Poisson random variables. We then describe a McMC method for obtaining an approximated sample from the collection when the individual random variables are all conditioned to be positive. Finally we describe the CFTP modification which allows us to obtain a perfect sample from the desired distribution.

We fix a given index set $S = \{1, \dots, k\}$ and assign a *rate* μ_A to each non-empty subset $A \subseteq S$. Consider a collection of *independent* Poisson random variables $\{Y_A : A \subseteq S\}$ indexed by these subsets, such that Y_A has mean μ_A, and define

$$X_i \quad = \quad \sum_{B : i \in B} Y_B \,.$$

Then $\{X_1, \dots, X_k\}$ is a *collection of correlated Poisson random variables*. We also define

$$X_A = \sum_{B:A\subseteq B} Y_B \, ,$$

and set

$$\lambda_A = \mathbb{E}[X_A] = \sum_{B:A\subseteq B} \mathbb{E}[Y_B] \, ,$$

abbreviating

$$\lambda_i = \lambda_{\{i\}} \, .$$

As described in §4, such collections arise naturally in the context of the Boolean model in stochastic geometry, when the index set $\{1, \ldots, k\}$ is replaced by a finite set $S = \{x_1, ..., x_k\}$ of *conditioning points* and Y_B counts the number of grains covering exactly the subset B of conditioning points and no others, while $X_{\{x_i\}}$ counts the number of grains covering the conditioning point x_i. In this context, motivated by problems from the oil industry, Lantuéjoul [12] has investigated simulation of such Boolean models conditioned to cover the conditioning points: the question of converting this to perfect simulation, as discussed in [11], leads directly to the question of how to conduct perfect simulation of correlated Poisson random variables conditioned to be positive.

2.1 McMC for collection of correlated Poisson random variables conditioned to be positive

To carry out McMC for $\{X_1, \ldots X_k\}$ we replace random variables by immigration-death processes. For each non-empty subset $A \subseteq S$ we define an immigration-death process Y_A of immigration rate μ_A and unit death rate, requiring the Y_A to be independent. By detailed balance Y_A has Poisson equilibrium of mean μ_A. Now we define X_A for any non-empty subset $A \subseteq S$ by

$$X_A(t) = \sum_{B:A\subseteq B} Y_B(t) \, . \tag{1}$$

Abbreviating $X_{\{i\}} = X_i$, it follows that the equilibrium distribution for $\{X_1, \ldots, X_k\}$ is a collection of correlated Poisson random variables. The McMC procedure is now straightforward: simulate the Y_A processes for a long time, then construct the X_i using Equation (1). Of course this is using a sledgehammer to crack a nut: there is no reason not to sample directly from the Y_A equilibria. However the McMC approach can easily be modified to allow for conditioning.

Considered alone, the process X_A is an immigration-death process of immigration rate λ_A and unit death rate. Moreover

$$\text{if } B \subseteq A \, , \text{ then } X_A \widetilde{\preceq} X_B \, , \tag{2}$$

where $X_A \widetilde{\preceq} X_B$ means that

- if there is an immigration in X_A at time t then there is an immigration in X_B at that time,

- and if there is a death in X_B at time t then there is a death in X_A at that time.

It is possible to establish a converse to the above as follows: given a collection of immigration-death processes $\{X_A : A \subseteq S\}$ satisfying the relation (2), and *co-adapted* in the sense that the rate-computations are based on conditioning on a common past, then independent immigration-death processes Y_A of immigration rate μ_A and unit death rate can be constructed such that Equation (1) applies, using the Möbius formula [1, chapter 5, page 53] to deduce

$$Y_B(t) \quad = \quad \sum_{A:B\subseteq A} (-1)^{\#(A\setminus B)} \; X_A(t)\,. \tag{3}$$

We now consider how to modify the above McMC procedure so as to produce correlated Poisson random variables conditioned to be positive. Since the underlying Markov chain is reversible, all that has to be done is to forbid transitions which lead to one or more of the X_i becoming zero. To do this in a way which maintains the coupling required for CFTP methods, we need to prescribe replacement death times for the potential death times which are thus prevented from becoming actual; accordingly we introduce independent unit-rate Poisson processes $\mathcal{Z}_A$, one for each non-empty subset $A \subseteq S$.

We define the immigration-death process $\widetilde{Y}_A$ to be based on the same immigration-death structure as Y_A, but forbidding deaths which would lead to one of the target random variables $X_1, \ldots, X_k$ becoming zero. Namely, if t is a death time of Y_A such that $X_i(t) = 0$ for some $i \in A$, then the death time t of Y_A will not be accepted by $\widetilde{Y}_A$, and we set $\widetilde{Y}_A(t) = \widetilde{Y}_A(t-)$ rather than $\widetilde{Y}_A(t-) - 1$. Furthermore, we use the next incident $\hat{t}$ of $\mathcal{Z}_A$ as a new potential death time of $\widetilde{Y}_A$. We call this *perpetuation* of $\widetilde{Y}$ from time t to $\hat{t}$. The interval $[t, \hat{t})$ is a *perpetuation interval*, and $\hat{t}$ is a *perpetuation time* of $\widetilde{Y}_A$. Summarizing, if $\widetilde{Y}_A(t-) = 1$, if Y_A experiences a death at t, and if there exists $i \in A$ such that

$$\widetilde{Y}_B(t) = 0 \qquad \text{for all } B \neq A \text{ such that } i \in B\,,$$

then perpetuation takes place. In particular, perpetuation increases $\widetilde{Y}_A$ by at most 1 compared with the original Y_A.

Since we use independent unit rate Poisson processes $\mathcal{Z}_A(t)$ $(A \subseteq S)$ to provide perpetuation times, $\widetilde{Y}_A$ $(A \subseteq S)$ inherits the unit death rate of Y_A subject to censorship of those transitions which would lead to $X_i(t) = 0$ for some $i \in A$.

Now if we define

$$\begin{aligned} \widetilde{X}_B(t) &= \sum_{A:B\subseteq A} \widetilde{Y}_A(t)\,, \\ \widetilde{X}_i(t) &= \widetilde{X}_{\{i\}}(t)\,, \end{aligned}$$

then detailed balance shows that the equilibrium distribution of

$$\widetilde{X} \quad = \quad \left(\widetilde{X}_1, \ldots, \widetilde{X}_k\right)$$

is that of a collection of correlated Poisson random variables conditioned to be positive.

The pseudo-code `McMC-CCP` displayed in Table 1 gives an explicit description of how we can derive a simulation of $\widetilde{X}$ over a time interval $[a, b]$ from a realization of $(Y_A : A \subseteq S)$ (equivalently $(X_A : A \subseteq S)$) and $(\mathcal{Z}_A : A \subseteq S)$ over the same time interval.

```
McMC-CCP([a,b]):
    initialize Ỹ_A = Y_A(a) for all A ⊆ S
    repeat
        for i = 1 to k repeat
            X̃_i = Σ_{A:i∈A} Ỹ_A
        A ← index of next incident
        if next incident is immigration then
            Ỹ_A ← Ỹ_A + 1
        if next incident is death then
            if X̃_j ≠ 0 for all j ∈ A then
                Ỹ_A ← Ỹ_A − 1
            else
                label next 𝒵_A-incident as death
        until no more incidents
    return (X̃_i : i = 1,...,k).
```

The algorithm McMC−CCP performs McMC to obtain an approximate sample of conditioned correlated Poisson random variables, using unconditioned immigration-death processes $\{Y_A\}$ and associated perpetuation Poisson processes $\{\mathcal{Z}_A\}$.

Table 1: McMC−CCP algorithm

2.2 Extended state-space CFTP in a general context

Dominated CFTP is presented in [10] in a general context. In this paper we require a further modest extension, *extended state-space CFTP*. Here we describe this in a general context, then describe how to use it to get perfect samples from a correlated Poisson random variables conditioned to be positive.

Our aim is to obtain a perfect sample from the equilibrium distribution π of a discrete or continuous time Markov chain $R = \{R(t), t \geq 0\}$ taking values in a state space $\mathfrak{X}$. The difficulties here are that the Markov chain R may not be uniformly ergodic, and even if a relevant partial order exists we do not wish to assume either monotonicity or the existence of maximal and minimal elements. The main idea to overcome this problem is: insert the original unordered state-space $\mathfrak{X}$ as a subset of a bigger ordered space $\widetilde{\mathfrak{X}}$, with an "upper/lower-sandwich" process living on $\widetilde{\mathfrak{X}}$ and eventually being absorbed in $\mathfrak{X}$. The evolution of the stochastic process $\widetilde{R}$ is chosen so that once it visits $\mathfrak{X}$ it follows the stochastic dynamics of R.

Specifically, we identify a sequence of majorizing processes

$$\widetilde{R}^{(n)} \quad = \quad \left\{\widetilde{R}^{(n)}(t) : \ -n \leq t \leq \infty\right\}, \qquad \text{such that } n = 1, 2, 4, \ldots,$$

which are identically distributed up to a shift in time, which live on the superset $\widetilde{\mathfrak{X}}$ of $\mathfrak{X}$ and satisfying the following conditions:

1. $\widetilde{\mathfrak{X}}$ possesses a partial order relation $\preceq$, and if $a \in \widetilde{\mathfrak{X}}$ and $b \in \mathfrak{X} \subset \widetilde{\mathfrak{X}}$ are such that $a \preceq b$ then $a = b$ (so $\mathfrak{X}$ is at the base of the partially ordered space $\widetilde{\mathfrak{X}}$);

2. $\widetilde{R}^{(n)}$ is eventually absorbed in the subspace $\mathfrak{X}$, so that

$$\mathbb{P}\left[\widetilde{R}^{(n)}(u) \in \mathfrak{X}\right] \quad \to \quad 1 \tag{4}$$

for any fixed u as $n \to \infty$;

3. $\widetilde{R}^{(n)}$ evolves according to the stochastic dynamics of R once it hits $\mathfrak{X}$;

4. For all $-m \leq -n \leq t \leq 0$

$$\widetilde{R}^{(m)}(t) \quad \preceq \quad \widetilde{R}^{(n)}(t)\,, \tag{5}$$

Then we have the following CFTP result.

Theorem 1 *Suppose that the Markov chain R takes values in a state space $\mathfrak{X}$, and the processes $\widetilde{R}^{(n)}$ live on the partially ordered superset $\widetilde{\mathfrak{X}}$ of $\mathfrak{X}$ and satisfy conditions 1–4. Suppose further that $R(t)$ converges weakly to a unique equilibrium distribution π as $t \to \infty$. Define*

$$T \quad = \quad \inf\left\{n \geq 1 : \widetilde{R}^{(n)}(0) \in \mathfrak{X}\right\}\,. \tag{6}$$

Then $T < \infty$ almost surely and $\widetilde{R}^{(T)}(0)$ follows the equilibrium distribution π.

Proof : It follows from Equation (4) in condition 2 above that $T < \infty$ almost surely. By Equation (5) of condition 4, if $n > T$ then

$$\widetilde{R}^{(n)}(0) \preceq \widetilde{R}^{(T)}(0)\,.$$

Consequently by condition 1 above $\widetilde{R}^{(n)}(0) = \widetilde{R}^{(T)}(0)$ for such n. Accordingly

$$\widetilde{R}^{(T)}(0) = \lim_{n\to\infty} \widetilde{R}^{(n)}(0)$$

exists almost surely. By condition 3, and the fact that the $\widetilde{R}^{(n)}$ are identically distributed up to time shifts, $\widetilde{R}^{(n)}(0)$ has the same equilibrium distribution as that of $R(n)$, so

$$\lim_{n\to\infty} \mathrm{Law}(\widetilde{R}^{(n)}(0)) \quad = \quad \lim_{n\to\infty} \mathrm{Law}(R(n))\,,$$

where R is started at time 0 with the common hitting distribution of the $\widetilde{R}^{(n)}$ on $\mathfrak{X}$. But $R(n)$ converges weakly to π. Therefore, $\widetilde{R}^{(T)}(0)$ has distribution π as required. $\square$

In the following subsection we describe how to construct $\widetilde{R}^{(n)}$ as an ordered triple of processes, and in the next section we show that conditions 1–4 hold, so that we can use $\widetilde{R}^{(n)}$ to get a perfect sample from the desired distribution.

2.3 CFTP for correlated Poisson random variables conditioned to be positive

We need to introduce some further notation in order to explain how we can use extended state-space CFTP to produce perfect samples of correlated Poisson random variables conditioned to be positive.

Definition 1 *Define $\mathfrak{X} \subset \widetilde{\mathfrak{X}}$ as follows*

$$\begin{aligned} \mathfrak{X} &= \{(y,v) : y, v \text{ are } n\text{-tuples of nonnegative integers}\} \, , \\ \widetilde{\mathfrak{X}} &= \Big\{(y, v^{min}, v^{max}) : y, v^{min}, v^{max} \text{ are } n\text{-tuples of nonnegative integers} \\ &\qquad\qquad \text{with } v^{min} \leq v^{max}\Big\} \, , \end{aligned}$$

where the n-tuples are all of length $2^k - 1$ indexed by the non-empty subsets of S, the inequality $v^{min} \leq v^{max}$ is interpreted element-wise, and in fact v, v^{min}, v^{max} only contain entries which are 0 or 1. The embedding $\mathfrak{X} \subset \widetilde{\mathfrak{X}}$ is given by

$$(y,v) \quad \mapsto \quad (y,v,v) \, .$$

The partial ordering on $\widetilde{\mathfrak{X}}$ is given by

$$(y^{(1)}, v^{(1),min}, v^{(1),max}) \quad \preceq \quad (y^{(2)}, v^{(2),min}, v^{(2),max})$$

exactly when $y^{(1)} = y^{(2)}$, $v^{(2),min} \leq v^{(1),min}$ and $v^{(1),max} \leq v^{(2),max}$, interpreting equalities and inequalities element-wise.

This ordering on $\widetilde{\mathfrak{X}}$ corresponds exactly to the inclusion ordering for order intervals which was used in set-valued CFTP and, implicitly, in dominated CFTP in earlier formulations [6, 7, 8].

It is apparent that this partial ordering and embedding of $\mathfrak{X}$ in $\widetilde{\mathfrak{X}}$ satisfies condition 1 preceding Theorem 1.

We now describe how to use the unconditional immigration-death processes (Y_A : non-empty $A \subseteq S$) (assumed to be in statistical equilibrium) and the perpetuating incident processes ($\mathcal{Z}_A$: non-empty $A \subseteq S$) to construct realizations of $\widetilde{R}^{(n)}$ over the time interval $[-n, 0]$. Note that we can use time-reversibility of the Y_A to extend their trajectories as far back into the past as may be required, fixing their values at time 0 to be samples from the unconditioned Poisson equilibrium of these processes. Extension of the $\mathcal{Z}_A$ is even easier!

For $-n \leq t \leq 0$ we define

$$\widetilde{R}^{(n)}(t) \quad = \quad \Big((Y_A(t), V_A^{(n),min}(t), V_A^{(n),max}(t)) : A \subseteq S\Big)$$

where Y_A is the *unconditional* immigration-death process assigned to the subset A, and the $V^{(n),min}$, $V^{(n),max}$ processes are defined together recursively in the following way.

Initially, at time $-n$, we set

$$\begin{aligned} V_A^{(n),min}(-n) &= 0 \\ V_A^{(n),max}(-n) &= 1 \end{aligned}$$

for all non-empty $A \subseteq S$. Updates of these processes occur only at death times of corresponding Y processes and perpetuation times of corresponding $\mathcal{Z}$ processes.

For convenience of notation we define

$$\begin{aligned} \widetilde{X}_i^{(n),max}(t) &= \sum_{B:i\in B} \Big\{Y_B(t) + V_B^{(n),max}(t)\Big\} \, , \\ \widetilde{X}_i^{(n),min}(t) &= \sum_{B:i\in B} \Big\{Y_B(t) + V_B^{(n),min}(t)\Big\} \, . \end{aligned} \tag{7}$$

Suppose Y_A experiences a death at time $t \in [-n, 0]$. Then $V_A^{(n),min}$, respectively $V_A^{(n),max}$, may be set equal to 1 when there is a danger of causing zeros in some of the opposite processes $\widetilde{X}^{(n),max}$, respectively $\widetilde{X}^{(n),min}$:

[*min-death*] If $\widetilde{X}_i^{(n),max}(t-) \leq 1$ for some $i \in A$, so that after the Y_A-death there would be a zero in some *upper* process $\widetilde{X}_i^{(n),max}$, then we set the *lower* component $V_A^{(n),min}(t) = 1$.

[*max-death*] If $\widetilde{X}_i^{(n),min}(t-) \leq 1$ for some $i \in A$, so that after the Y_A-death there would be a zero in some *lower* process $\widetilde{X}_i^{(n),min}$, then we set the *upper* component $V_A^{(n),max}(t) = 1$.

On the other hand suppose that $\mathcal{Z}_A$ experiences a perpetuation incident at time t. Then we can set $V_A^{(n),min}$, respectively $V_A^{(n),max}$, equal to 0 when there is no danger of causing zeros in any of the opposite processes $\widetilde{X}^{(n),max}$, respectively $\widetilde{X}^{(n),min}$:

[*min-perpetuation*] If $\widetilde{X}_i^{(n),max}(t-) > 1$ for *all* $i \in A$, so that there is no danger of zeros in any *upper* process $\widetilde{X}_i^{(n),max}$, then we set the *lower* component $V_A^{(n),min}(t) = 0$.

[*max-perpetuation*] We need to consider the following two cases.

- $V_A^{(n),min}(t-) = 1$. In this case, if $\widetilde{X}_i^{(n),min}(t-) > 1$ for *all* $i \in A$, then we set the *upper* component $V_A^{(n),max}(t) = 0$.
- $V_A^{(n),min}(t-) = 0$ and $V_A^{(n),max}(t-) = 1$. In this case, if $\widetilde{X}_i^{(n),min}(t-) \geq 1$ for *all* $i \in A$, then we set the *upper* component $V_A^{(n),max}(t) = 0$.

It remains to establish the conditions 1–4 required by Theorem 1, and we shall do that in the next section. First, we give a brief description of our CFTP algorithm. It begins by testing for whether $\widetilde{R}^{(n)}(0)$ lies in $\mathfrak{X}$, which is to say, whether $V_A^{(n),min}(0) = V_A^{(n),max}(0)$ for all non-empty $A \subseteq S$. If so then the vector of common values $X_i^{(n),min}(0) = X_i^{(n),max}(0)$ for $i = 1, \ldots, k$ is returned as a sample of the corresponding vector of correlated Poisson random variables conditioned to be positive. If there exist $A \subseteq S$ such that $V_A^{(n),min}(0) \neq V_A^{(n),max}(0)$ then we must choose a new starting time $-m < -n$ and repeat the simulation replacing n by m (as suggested in [13], in this case a sensible choice is to double n). Note that we must insist on extending backwards in such a way that we re-use the previous realization of $Y_A|_{[-n,0]}$ and $\mathcal{Z}_A|_{[-n,0]}$ for non-empty $A \subseteq S$: as we noted above, the reversibility arising from detailed balance is of technical assistance here. We proceed to extend backwards in time until minimal and maximal processes coalesce, and then return the common value of the coalesced processes at time 0. Table 2 gives pseudo-code summarizing the resulting algorithm `CFTP-CCP`.

3 Properties of the CFTP algorithm

To prove that our CFTP algorithm `CFTP-CCP` will indeed provide us with perfect samples from the desired distribution, we must establish a sequence of lemmas relating to the conditions 1–4.

```
CFTP-CCP(−n):
    extend Y, 𝒵 to [−n,0]
    if R̃^(n)(0) ∈ 𝔛 then
        for i = 1 to k repeat
            X̃_i(0) = Σ_{A:i∈A} (Y_A(0) + V_A^{(n),min}(0))
        return (X̃_i(0) : i = 1, ... , k)
    else
        CFTP-CCP(−2n)
```

The CFTP-CCP algorithm performs extended state-space CFTP to obtain an exact sample of conditioned correlated Poisson random variables.

Table 2: CFTP-CCP algorithm

First we must make sure that condition 1 applies, which is to say, the process $\widetilde{R}^{(n)}$ genuinely stays in the state-space $\widetilde{\mathfrak{X}}$.

Lemma 1 *For $V^{(n),min}$, $V^{(n),max}$ defined as above we have*

$$V_A^{(n),min}(t) \quad \leq \quad V_A^{(n),max}(t) \tag{8}$$

for all non-empty $A \subseteq S$ and for $-n \leq t \leq 0$. In particular, the state-space for $\widetilde{R}^{(n)}$ can be taken to be $\widetilde{\mathfrak{X}}$.

Proof : First note that by construction Equation (8) must hold at time $-n$.

Now suppose that Equation (8) holds up to, but possibly not including, time t. Therefore we have

$$\widetilde{X}_i^{(n),min}(t-) \quad \leq \quad \widetilde{X}_i^{(n),max}(t-) \tag{9}$$

for all $i \in S$.

It suffices to consider t such that either Y_A experiences a death at that time, or $\mathcal{Z}_A$ delivers a perpetuation incident. In the case of a death the relationship (9) and the crossover between [*min-death*] and [*max-death*] assure us that $V^{(n),min}$ can be set to 1 only if $V^{(n),max}$ is, so Equation (8) must also hold at time t.

The case of a perpetuation incident is entirely similar: If the perpetuation happens in both upper and lower processes then $V^{(n),max}$ can be set to 0 only if $V^{(n),min}$ is, so Equation (8) must also hold at time t. If the perpetuation happens only in the upper process, we already have $V^{(n),min}(t) = 0$. So (8) holds no matter what is the value of $V^{(n),max}(t)$. The result follows by induction. □

Setting aside condition 2 for the moment, we turn to condition 3.

Lemma 2 *If $V^{(n),min} = V^{(n),max}$ at time $t > -n$ then this equality persists over all $[t, 0]$ and over that time interval the coupled process $Y + V^{(n),min} = Y + V^{(n),max}$ evolves as a copy of $\widetilde{Y}$.*

Proof : This is simply a matter of observing that if $V^{(n),min}$, $V^{(n),max}$ coincide at time t then so do $\widetilde{X}^{(n),min}$, $\widetilde{X}^{(n),max}$, and so there is no longer any distinction between min and max in the conditions [*min-death*], [*max-death*] for death and [*min-perpetuation*], [*max-perpetuation*] for perpetuation above, note that the case of perpetuation in the upper process only will not happen any more, while birth decisions are already identical. Consequently from time t onwards the death and perpetuation decisions are the same for $V^{(n),min}$, $V^{(n),max}$ so they must continue to coincide. Furthermore the effect of the death and perpetuation decisions on $Y + V^{(n),min}$, $Y + V^{(n),max}$ agrees with the prescription for the evolution of $\widetilde{Y}$ in our description of the `McMC-CCP` Algorithm in Table 1. □

Now consider condition 4. This follows if we can establish the following "funnelling" result:

Lemma 3 *For all $-n \leq -m \leq t \leq 0$ and all non-empty $A \subseteq S$ we have*

$$V_A^{(m),min}(t) \quad \leq \quad V_A^{(n),min}(t) \quad \leq \quad V_A^{(n),max}(t) \quad \leq \quad V_A^{(m),max}(t)\,.$$

Proof : The initial values $V^{(m),min}(-m)$, $V^{(m),max}(-m)$ are respectively minimal and maximal, so this funnelling certainly holds at time $t = -m$. Consider therefore an incident at time $t > -m$, and suppose the funnelling relationship holds at time $t-$. This carries through into a relationship

$$\widetilde{X}^{(m),min}(t-) \quad \leq \quad \widetilde{X}^{(n),min}(t-) \quad \leq \quad \widetilde{X}^{(n),max}(t-) \quad \leq \quad \widetilde{X}^{(m),max}(t-)\,.$$

If the incident is an immigration then all of $V^{(m),min}$, $V^{(m),max}$, $V^{(n),min}$ and $V^{(n),max}$ are not affected, and the relationship persists.

If the incident is a death then examination of [*min-death*] shows that $V^{(m),min}$ can be set to 1 only if $V^{(n),min}$ is set to 1, since $\widetilde{X}^{(n),max}(t-) \leq \widetilde{X}^{(m),max}(t-)$, and similar arguments using [*max-death*] apply to show $V^{(n),max}$ can be set to 1 only if $V^{(m),max}$ is set to 1.

If the incident is a perpetuation in both upper and lower processes then similar arguments concerning [*min-perpetuation*] and [*max-perpetuation*] show that $V^{(n),min}$ can be set to 0 only if $V^{(m),min}$ is set to 0, $V^{(m),max}$ can be set to 0 only if $V^{(n),max}$ is set to 0.

Now we consider the case, where the incident is a perpetuation only in the upper process. In this case, we have

$$V_A^{(m),min}(t-) \quad = \quad V_A^{(m),min}(t) \quad = \quad 0\,,$$

and

$$V_A^{(m),max}(t-) \quad = \quad 1\,.$$

Then it follows from Lemma 1 that

$$V_A^{(m),min}(t) \quad \leq \quad V_A^{(n),min}(t) \quad \leq \quad V_A^{(n),max}(t)\,.$$

So we only need to prove

$$V_A^{(n),max}(t) \quad \leq \quad V_A^{(m),max}(t)\,. \tag{10}$$

- If there is $i \in A$ such that $\widetilde{X}_i^{(m),min}(t-) = 0$ then according to [*max-perpetuation*], $V_A^{(m),max}(t)$ can not be set to 0. So $V_A^{(m),max}(t) = 1$, hence (10) holds.
- If $\widetilde{X}_i^{(m),min}(t-) \geq 1$ for all $i \in A$, then we set $V_A^{(m),max}(t) = 0$, and we have

  $$\widetilde{X}_i^{(n),min}(t-) \quad \geq \quad \widetilde{X}_i^{(m),min}(t-) \quad \geq \quad 1\,.$$

  for all $i \in A$. There are two cases to be considered:

  - $V_A^{(n),min}(t-) = 1$
    In this case, since $V_A^{(m),min}(t-) = 0$, we actually know

    $$\widetilde{X}_i^{(n),min}(t-) \quad > \quad \widetilde{X}_i^{(m),min}(t-)$$

    and hence also $\widetilde{X}_i^{(n),max}(t-) > 1$ for all $i \in A$. Then it follows from [*max-perpetuation*] and [*min-perpetuation*] conditions that we have

    $$V_A^{(n),min}(t) \quad = \quad V_A^{(n),max}(t) \quad = \quad 0\,.$$

  - $V_A^{(n),min}(t-) = 0$
    In this case, we have either $V_A^{(n),max}(t-) = 1$ or $V_A^{(n),max}(t-) = 0$.
    * If $V_A^{(n),max}(t-) = 1$ then, since

      $$\widetilde{X}_i^{(n),min}(t-) \quad \geq \quad \widetilde{X}_i^{(m),min}(t-) \quad \geq \quad 1\,,$$

      it follows from [*max-perpetuation*] that $V_A^{(n),max}(t) = 0$.
    * If $V_A^{(n),max}(t-) = 0$, then $V_A^{(n),max}(t) = 0$ is automatic since t is a perpetuation time not a death time.

  So in any case, (10) holds.

This means the funnelling relationship persists to time t if it has been maintained up to time $t-$. The result follows by induction along the sequence of incidents over the whole time interval $[-m, 0]$. □

Finally we turn to condition 2. This is a consequence of the following lemma.

Lemma 4 *For all sufficiently large n, the processes $V^{(n),min}$, $V^{(n),max}$ coalesce at some time in the interval $[-n, 0]$.*

Proof : The lemma follows if there are times at which the upper processes $V_A^{(n),max}$ all vanish. For by construction it then follows coalescence occurs at that time.

Fix $-T < t < 0$, $\delta > 0$ and for $i = 1, \cdots, k$ set

$$D_i = \left[\widetilde{X}_i^{(T),max} \text{ experiences an immigration in the time interval } [t,\ t+\delta]\right]\,.$$

Then the D_i are increasing events based on independent events of the form that Y_A has an immigration in $[t, t+\delta]$. So the FKG inequality [5, Theorem 2.4] applies and consequently

$$\mathbb{P}\left[\bigcap_{i=1}^k D_i\right] \quad \geq \quad \prod_{i=1}^k \mathbb{P}\left[D_i\right]\,.$$

On the other hand there is $\rho > 0$ such that $\mathbb{P}[D_i] > \rho$ for all i. It follows that $\prod_{i=1}^k \mathbb{P}[D_i] > 0$. (An alternative line of reasoning constructs an estimate based on all of the Y_A experiencing immigration, thus dispensing with the FKG inequality.)

Now suppose that at time $t + \delta$ the event $\bigcap_{i=1}^k D_i$ occurs. Let $\mathcal{S}_t$ be the family of subsets $A \subseteq S$ such that Y_A experiences at least one immigration in $[t, t + \delta]$. If $\bigcap_{i=1}^k D_i$ occurs then the union $\bigcup_{A \in \mathcal{S}_t} A$ must cover S, moreover clearly $\mathcal{S}_t$ has no more than $2^k - 1$ members.

Now there are at most $2^k - 1$ perpetuated values among $V_A^{(T),max}$, for all $A \subseteq S$. Coalescence of $V_A^{(T),min}$ and $V_A^{(T),max}$ will occur if, for all A yielding perpetuated values, *both* the Poisson process $\mathcal{Z}_A$ experiences incidents in the time interval $[t+\delta, t+2\delta]$ and *also* the Y_A does not hit zero in the same time interval, since then all the $V_A^{(T),max}$ will be zero. But a simple stochastic comparison, using the Strong Markov property and the representation of immigration-death processes in terms of Poisson inputs and exponential death times, shows that this happens with a conditional probability (conditioned on $\bigcap_i D_i$) greater than

$$(1 - \mathbb{P}[\text{Poisson}(\delta) = 0])^{(2^k-1)} \times \mathbb{P}\left[\text{Exponential}((2^k - 1)) > 2\delta\right] \;=\; \rho_1 > 0\,. \tag{11}$$

Let C_t be the event that $\bigcap_i D_i$ occurs together with the event whose probability is given in Equation (11) above. Then

$$\mathbb{P}[C_t] \quad \geq \quad (\rho)^k \times \rho_1 \quad > \quad 0\,.$$

However C_t is actually independent of events in the past of time t. We may therefore apply the second Borel-Cantelli lemma to show that, almost surely, infinitely many of the $C_{2r\delta}$ occur (for $r = 1, 2, \ldots$). It follows that coalescence is almost sure to happen. □

The following result is an immediate consequence of the above lemmas and Theorem 1 and the existence of a unique limiting equilibrium for processes $\widetilde{X}$ in §2.

Theorem 2 *The algorithm described in* §2 *almost surely terminates in finite time, moreover the output of the algorithm has the distribution of the corresponding array of correlated Poisson random variables conditioned to be positive.*

The above work shows that our CFTP algorithm does indeed produce a perfect sample of the corresponding collection of correlated Poisson random variables conditioned to be positive.

4 Application to simulation of the conditional Boolean model

In this section, we show how our CFTP algorithm can be used to obtain a perfect sample from a conditional Boolean model.

4.1 Relationship between the conditional Boolean model and correlated Poisson random variables conditioned to be positive

Recall that in section 2, we introduced the distribution of correlated Poisson random variables $(X_1, \ldots, X_k)$ and showed it could be obtained as the equilibrium distribution of the process $(X_1(t), \ldots, X_k(t))$, where

$$X_i(t) = \sum_{A:i\in A} Y_A(t), \qquad \text{for } i = 1, \ldots, k\,,$$

for independent immigration-death processes $Y_A(t)$ $(A \subseteq S)$.

Now if we consider a Boolean model as an equilibrium distribution for a spatial immigration-death process then with appropriate assignation of conditioning points $x_1, \ldots, x_k$, and rates μ_A, the process $X_B(t)$ can be interpreted as the time evolving number of germ-grain pairs whose grains cover all the conditioning points in B, while the process $Y_A(t)$ describes the number of germ-grain pairs whose grains cover all conditioning points in A, but no conditioning point which is not in A. For example, if $S = \{x_1, x_2\}$, while $B = \{x_1\}$, then $X_B(t)$ counts the number of germ-grain pairs whose grains cover the conditioning point x_1. We divide this process $X_B(t)$ into processes $Y_{\{x_1\}}(t)$ and $Y_{\{x_1,x_2\}}(t)$, where $Y_{\{x_1\}}(t)$ counts the number of germ-grain pairs whose grains cover x_1, but not x_2, and $Y_{\{x_1,x_2\}}(t)$ counts the number of germ-grain pairs whose grains cover both x_1 and x_2. If we convert the Boolean model into a Poisson process on a space marked using the grains, then the space of germs leading to grains covering x_1 is divided into two disjoint subregions; $Y_{\{x_1\}}(t)$ counts the number of germs which fall in one of the subregions, and $Y_{\{x_1,x_2\}}(t)$ counts the other. Thus the number of germ-grain pairs in the subregion associated with A is an immigration-death process $Y_A(t)$, independent of other such immigration-death processes until the conditioning is applied by means of forbidding certain transitions.

We have noted that some of the processes $Y_A(t)$ might have zero immigration rate. In the conditional Boolean model, this corresponds to a set A of conditioning points which can not simultaneously be covered by a single germ-grain pair. In the case of grains which are disks of constant radius, there will then be at least two points in A which are separated by a distance of greater than $2r$, where r is the radius of a typical grain.

From the above it can be seen how we can use perfect samples from conditional correlated Poisson random variables to produce perfect samples from a conditional Boolean model.

4.2 Perfect simulation for the conditional Boolean model

For the sake of simplicity we consider only the case of disks of constant radius, though the method can be extended to cover the case when the typical grain is random but of bounded diameter.

Perfect simulation follows once one establishes how to simulate sorted incident times of all $Y_A(t)$ and $\mathcal{Z}_A(t)$ over $[-T, 0]$. In essence all one has to do is to simulate a stream of candidate incident times of time-varying intensity, then for each incident time choose whether it is a $\mathcal{Z}$-type or a Y-type incident using appropriate probabilities.

The time-varying rate of the underlying stream of candidate incidents should be taken to be

$$N + \left(\sum_A Y_A\right) + \lambda \times \text{area}(W)\,,$$

where N is the total number of processes $\mathcal{Z}_A$ with $\mu_A \neq 0$, and assignations of type ($\mathcal{Z}$, Y-death, Y-birth) should be made so as to deliver corresponding rates

$$N\,, \quad \sum_A Y_A\,, \quad \lambda \times \text{area}(W)\,.$$

For $\mathcal{Z}$-type incidents, one assigns the type (indicating to which of the various $\mathcal{Z}_A$ processes the incident belongs) uniformly at random. For Y-type death incidents one assigns the type (indicating to which of the various Y_A processes the incident belongs) randomly but weighted by the sizes of the corresponding Y_A. For Y-type birth incidents, one chooses a location uniformly at random from the bounded window W and then assigns the incident to Y_A if this location falls in the region

$$E_A \quad = \quad \left(\bigcap_{x\in A} x \oplus \check{G}\right) \setminus \left(\bigcup_{y\notin A} y \oplus \check{G}\right) \tag{12}$$

where $\check{G} = -G = \{-x:\ x \in G\}$ and G is a typical grain (since then the corresponding disk will cover exactly those conditioning points constituting the indexing set A). Birth incidents can be discarded if their locations do not fall in any of these regions.

Perfect simulation is then conducted according to the algorithm `CFTP-CCP` as given in Table 2. At the conclusion of the algorithm the perfectly sampled conditional Boolean model is built up using the locations of surviving birth incidents, supplemented by a sample from the *unconditional* Boolean model on the region

$$W \setminus \left(\bigcup_{y\in S} y \oplus \check{G}\right)\,,$$

which by construction covers no conditioning points. (This can be done easily by thinning the unconditional Boolean model on the full region W.)

In the following, we consider plots of conditional Boolean models obtained from the above perfect simulation method. Clearly, given the configuration of conditioning points and the radius of a typical grain, the results will be affected by the value of the intensity. We give plots for simulations using varying patterns of conditioning points (random pattern, oblique lattice, square lattice) and medium and high intensities of the underlying Poisson germ process. For low intensity (less than 30 germs per unit area) and closely-spaced conditioning points (such that a grain can cover more than one conditioning point) the algorithm failed to coalesce within the memory constraints of our computer (Sun Ultra 1, 64 MBytes RAM); the maximum and minimum processes then take a long time to come to agreement, since initially the perpetuated maximum grains disappear only when minimum grains cover conditioning points (itself a rare event), and moreover perpetuated minimum grains tend to disappear quickly back to zero since conditioning points are typically covered by several perpetuated maximum grains.

Our plots are based on intensities of 30 (medium) and 45 (high) germs per unit area. The same general features are displayed in each case: there are many extra grains induced by the requirement to cover all conditioning points. For high intensity the resulting figure differs less from the unconditional case, since much of the window is covered anyway, but there is still an excess of grains covering conditioning points. In each case we compute the number of grains covering the conditioning points and the excess of this over the number to be expected for a comparable simulation of an unconditional Boolean model. The excess is much larger in the case of a random pattern of conditioning points; this is reasonable since a randomly scattered pattern affects a much greater area of the window. The excess does not appear to depend strongly on the underlying germ intensity. This comes as no surprise. The excess, when generated by our algorithm, is merely the sum of all perpetuated grains: $\sum V_A$ in the notation of section 2. Since the V_A are only over 0 or 1, the scope for variation of the total excess is limited.

Note that in all the figures conditioning points are marked as "+".

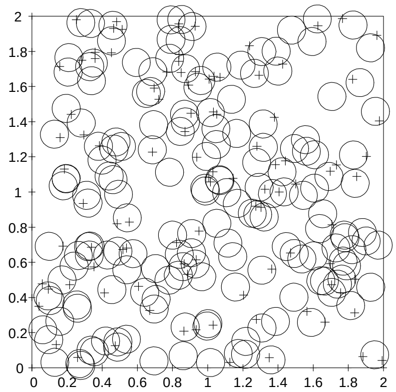

Perfect sample from a conditional Boolean model with
- $k = 100$ conditioning points placed randomly in a window W of width $\ell = 2$ units.
- Intensity of underlying germ process $\lambda = 30$.
- Radius of typical grain $r = 0.08$.

Simulation yields $n = 173$ grains, of which 103 cover at least one conditioning point. The unconditional model would yield a Poisson number of grains, mean number 120, of which we would expect 47.41 to cover at least one of the $k = 100$ conditioning points. So the excess number of grains from the conditional Boolean model covering the conditioning points is 55.59.

Figure 1: Conditional Boolean model, random pattern

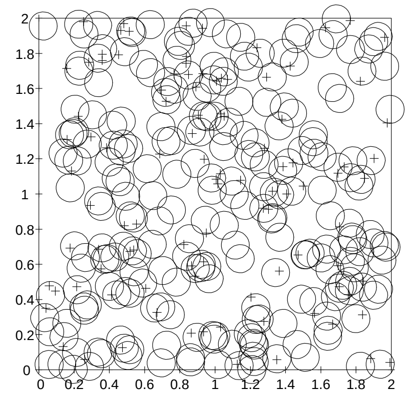

Perfect sample from a conditional Boolean model with
- $k = 100$ conditioning points placed randomly in a window W of width $\ell = 2$ units.
- Intensity of underlying germ process $\lambda = 45$.
- Radius of typical grain $r = 0.08$.

Simulation yields $n = 246$ grains, of which 131 cover at least one conditioning point. The unconditional model would yield a Poisson number of grains, mean number 180, of which we would expect 71.11 to cover at least one of the $k = 100$ conditioning points. So the excess number of grains from the conditional Boolean model covering the conditioning points is 59.89.

Figure 2: Conditional Boolean model, random pattern

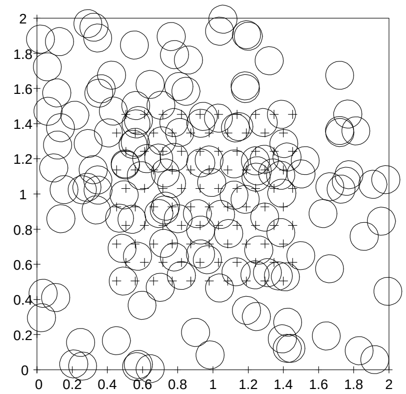

Perfect sample from a conditional Boolean model with
- $k = 100$ conditioning points placed on a lattice based on a parallelogram of horizontal side length $2/19$, height $2/19$ and angle 63.43^o in a window W of width $\ell = 2$ units.
- Intensity of underlying germ process $\lambda = 30$.
- Radius of the typical grain $r = 0.08$.

Simulation yields $n = 149$ grains, of which 71 cover at least one conditioning point. The unconditional model would yield a Poisson number of grains, mean number 120, of which we would expect 47.41 to cover at least one of the $k = 100$ conditioning points. So the excess number of grains from the conditional Boolean model covering the conditioning points is 23.59.

Figure 3: Conditional Boolean model, parallelogram grid

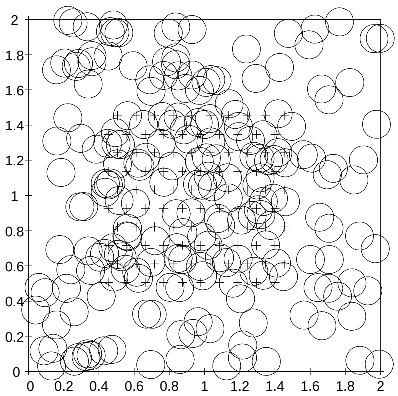

Perfect sample from a conditional Boolean model with
- $k = 100$ conditioning points placed on a lattice based on a parallelogram of horizontal side length $2/19$, height $2/19$ and angle 63.43^o in a window W of width $\ell = 2$ units.
- Intensity of underlying germ process $\lambda = 45$.
- Radius of the typical grain $r = 0.08$.

Simulation yields $n = 207$ grains, of which 97 cover at least one conditioning point. The unconditional model would yield a Poisson number of grains, mean number 180, of which we would expect 71.11 to cover at least one of the $k = 100$ conditioning points. So the excess number of grains from the conditional Boolean model covering the conditioning points is 25.89.

Figure 4: Conditional Boolean model, parallelogram grid

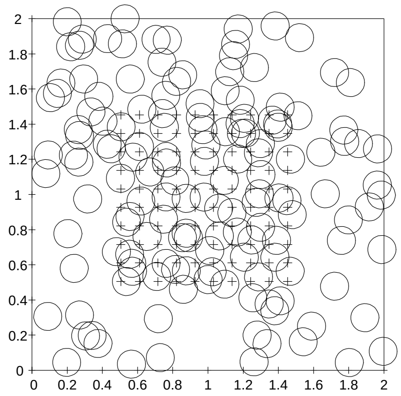

Perfect sample from a conditional Boolean model with
- $k = 100$ conditioning points placed on a square lattice of side length $2/19$ in a window W of width $\ell = 2$ units.
- Intensity of underlying germ process $\lambda = 30$.
- Radius of the typical grain $r = 0.08$.

Simulation yields $n = 150$ grains, of which 74 cover at least one conditioning point. The unconditional model would yield a Poisson number of grains, mean number 120, of which we would expect 47.41 to cover at least one of the $k = 100$ conditioning points. So the excess number of grains from the conditional Boolean model covering the conditioning points is 26.59.

Figure 5: Conditional Boolean model, square grid

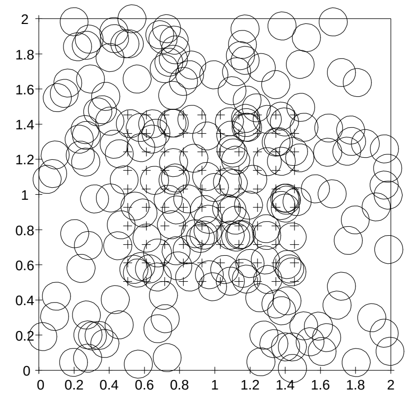

Perfect sample from a conditional Boolean model with
- $k = 100$ conditioning points placed on a square lattice of side length $2/19$ in a window W of width $\ell = 2$ units.
- Intensity of underlying germ process $\lambda = 45$.
- Radius of the typical grain $r = 0.08$.

Simulation yields $n = 212$ grains, of which 98 cover at least one conditioning point. The unconditional model would yield a Poisson number of grains, mean number 180, of which we would expect 71.11 to cover at least one of the $k = 100$ conditioning points. So the excess number of grains from the conditional Boolean model covering the conditioning points is 26.89.

Figure 6: Conditional Boolean model, square grid

5 Checking for correctness

Implementations of simulation methods should always be checked out carefully, to protect against mistakes in coding as well as to confirm the theory. In this section we summarize several simple examples which provided checks for the correctness of the method. The programs are implemented in `C`. The simulation time interval is initialized at $[-2, 0]$ in each case, though of course the algorithm can extend this further back in time if required. Window W is a unit square. G is set to be a disk of radius r, λ is the intensity of the germ process. To check correctness, in each case we collect 10000 samples and then carry out a χ^2 test using appropriate grouping. In the following we only report summarized results. Full details of implementation, results, and derivations of the exact equilibrium distribution functions for the examples can be found in the associated research report [3], available as Preprint 350 at

`http://www.warwick.ac.uk/statsdept/Staff/WSK/ppt.html`.

1. Single conditioning points.
   $S = \{x_1\},\ x_1 = (0.5, 0.5),\ r = 0.5,\ \lambda = 5$. The resulting χ^2_5 value was 4.804.

2. Two separated conditioning points.
   $S = \{x_1, x_2\},\ x_1 = (0.25, 0.25),\ x_2 = (0.75, 0.25),\ r = 0.25,\ \lambda = 5$. The resulting χ^2_8 value was 9.368.

3. Two close conditioning points.
   $S = \{x_1, x_2\},\ x_1 = (0.25, 0.25),\ x_2 = (0.4, 0.25),\ r = 0.25,\ \lambda = 5$. The resulting χ^2_{15} value was 12.344.

4. Three close conditioning points in 3-fold symmetry.
   $S = \{x_1, x_2, x_3\}$. Place the three points in a configuration of 3-fold symmetry. $r = 0.3,\ \lambda = 8$. The resulting χ^2_{26} value was 19.993.

None of the above χ^2 values are significant.

6 Investigating algorithm performance

In this section we carry out some experiments to explore the way in which the run-time of the perfect simulation algorithm depends on various parameters.

It follows from the analysis in Section 4.2 that each unconditioned process Y_A is a spatial immigration-death process of *constant* immigration rate λ and unit death rate on E_A, where $A \subseteq S$ and E_A is given by Equation (12).

Recall that k is the number of conditioning points and r is the radius of a typical grain.

In this section we will investigate the dependence of T_c on λ, r, k. Note that throughout this section, we use the actual coalescence time T_c, not the time produced by doubling during the performance of Algorithm `CFTP-CCP`.

Experiment 1: dependence of coalescence time on grain radius

Holding fixed the intensity of the germ process and the number of conditioning points, we investigate the relation between coalescence time T_c and the radius r of a typical grain.

We place a set $S = \{(u_i, v_j) : \ i, j = 1, 2\}$ of $k = 4$ conditioning points on the vertices of a square of side-length 0.2 in a unit-width window W, and obtain independent samples of size 200 each for 99 equally spaced values of r ranging from 0.01 to 0.5. Figure 7 plots log(sample mean) of coalescence time versus r (the dotted line), and includes two theoretically derived asymptotics (the curved solid line for small radius, and the horizontal solid line for large radius). A discontinuity in the sample curve signals the effect of the possibility of single grains covering several conditional points.

The theoretical asymptotics are derived as follows.

First consider the case of grains so small that any one grain can cover at most one conditioning point ($r < 0.1$). Then the only relevant Y_A are the k processes for which $A = \{i\}$ is a singleton, and they are independent of each other. Coalescence occurs exactly when each of the $Y_{\{i\}}$ has experienced a death at a time when there is no perpetuation in $V_{\{i\}}^{max}$. For an individual $Y_{\{i\}}$ the first such event occurs approximately at the first instance of a Poisson process of rate $\lambda \pi r^2$ which is independently thinned with deletion probability

$$\int_0^\infty \exp(-t)\exp(-t)\mathrm{d}t \quad = \quad 1/2$$

Using this approximation, the thinned process is also a Poisson process of rate $\lambda\pi r^2/2$. Therefore

$$\mathbb{P}\left[T_c \leq t\right] \quad \approx \quad \left(1 - \exp(-\lambda\pi r^2 t/2)\right)^k \tag{13}$$

At the alternative extreme of large r ($r >> 0.1$), the conditioning points are likely to be covered by just one disk of the unconditional Boolean model: coverage probability for just one conditioning point on its own at time zero is

$$1 - \exp(-\lambda\pi r^2)\,,$$

and the coverage events are positively correlated. Accordingly the coalescence time T_c will be approximated by the maximum of the first perpetuation times for each of the N Y_A processes, so

$$\mathbb{P}\left[T_c \leq t\right] \quad \approx \quad \left(1 - \exp(-t)\right)^N \tag{14}$$

not depending on grain radius r.

We deduce the following theoretical asymptotics for the mean coalescence time:

$$\mathbb{E}\left[T_c\right] \quad \approx \quad \frac{2k}{\lambda\pi r^2}\sum_{i=0}^{k-1}\left(\begin{array}{c} i \\ k-1 \end{array}\right)(-1)^{k+1-i}\frac{1}{(k-i)^2} \tag{15}$$

for small r, and

$$\mathbb{E}\left[T_c\right] \quad \approx \quad N\sum_{i=0}^{N-1}\left(\begin{array}{c} i \\ N-1 \end{array}\right)(-1)^{N+1-i}\frac{1}{(N-i)^2} \tag{16}$$

for large r.

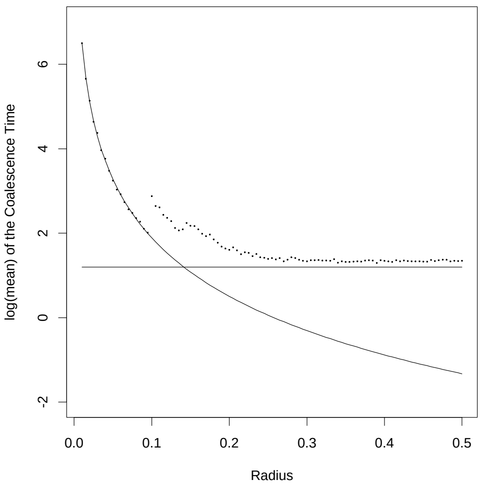

Figure 7: Graph of the $\log(\text{sample mean})$ of T_c versus grain radius r for fixed conditioning points $k = 4$ and value of $\lambda = 20$. The reason for the big jump at $r = 0.1$ is as follows. When $r = 0.1$, it is possible for one grain to cover more than one conditioning point. The jump shows the cost we pay for the consequent interaction between $\widetilde{Y}$ processes.

Experiment 2: dependence of coalescence time on intensity

Holding fixed the radius of a typical grain and the number of conditioning points, we now investigate the relation between coalescence time T_c and the intensity λ of the germ process.

We place a set $S = \{(u_i, v_j) : \ i, j = 1, 2, 3, 4\}$ of $k = 16$ conditioning points on the vertices of a square of side-length 0.2 in a unit-width window W, and obtain independent samples of size 200 each for 75 equally spaced values of λ ranging from 2 to 150. Figure 8 plots the log(sample mean) of coalescence time versus λ (dotted line) and again includes two theoretically derived asymptotics (solid lines).

The theoretical asymptotics are derived essentially as for Experiment 1, as follows.

First consider the case where λ is very large. Then all the conditioning points are likely to be covered by a single disk of the unconditional Boolean model so, with the same arguments as those in Experiment 1, the coalescence time T_c can also be approximated by the maximum of the N first perpetuation times for the Y_A processes. Hence the mean is approximated by asymptotic (16), which is independent of the value of λ.

On the other hand, for small values of λ, the coalescence time T_c will depend on the the value of r. In this experiment, we take $r = 0.08$, so it follows from the arguments in Experiment 1 that when λ is not very large, the distribution of T_c is approximately given by asymptotic (13). Hence the approximated theoretical mean is given by (15), but viewed as a function of λ.

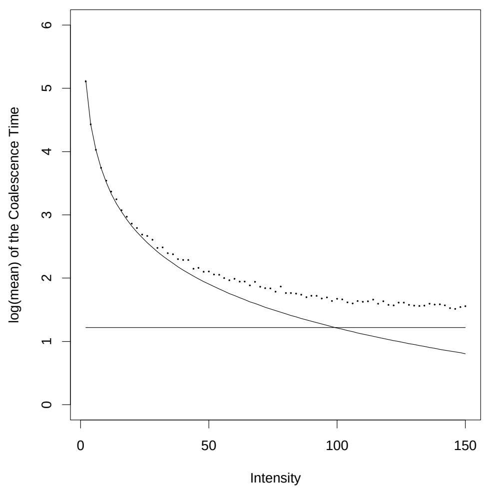

Figure 8: Graph of the $\log(\text{sample mean})$ of T_c versus intensity λ for fixed conditioning points $k = 16$ and radius of the grain $r = 0.08$.

Experiment 3: dependence of coalescence time on the number of conditioning points

Holding fixed the intensity of the germ process and the radius r of a typical grain, we investigate the relation between coalescence time T_c and the number of conditioning points.

On a unit-width window W, when k is small, it is likely that a single grain can cover at most one conditioning point as long as r and λ are not very large. Therefore, for small k, the distribution of T_c can also be approximated by equation (13), hence the mean is given by asymptotic (15) but now viewed as a function of k.

On the other hand, when k is very large, any disk of the unconditional Boolean model are likely to cover at least one conditioning point. So the distribution of T_c in this case can be approximated by equation (14), where again, N is the number of Y_A processes with $A \subseteq S$, $\mu_A \neq 0$. Hence the mean is given by asymptotic (16), but now viewed as a function of k.

We illustrate the relation between T_c and k when r is small. The experiment is designed as follows.

We take $r = 0.04,\ \lambda = 30$ and place k $(k = 1, \cdots, 50)$ conditioning points on the vertices of rectangles in a unit-width window W, each rectangular has width $1/11$ and length $1/6$. We obtain independent samples of size 200 for each k. By above settings, any grain can cover at most one conditioning point. Therefore, in this case, we would expect that the sample mean versus k should agree with the approximated theoretical mean curve given by asymptotic (15), as is confirmed by the plot of sample mean (dotted lines) and the two asymptotics of the theoretical mean versus number of conditioning points k in Figure 9.

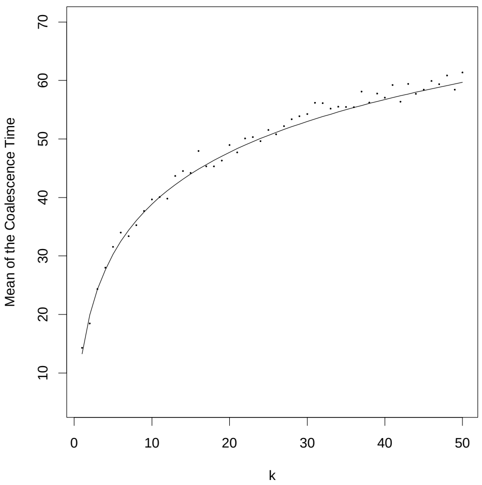

Figure 9: Graph of the sample mean of T_c versus k for fixed value of $\lambda = 30$ and radius of the grain $r = 0.04$.

Experiment 4: dependence of actual duration of algorithm on intensity

Finally we investigate the relation between actual duration ("wall-clock time") of the algorithm and λ, r and k. In each experiment, we recorded 200 independent actual duration times for collecting one sample for each value of r, λ or k and plotted the average duration time and 95% confidence limits about the mean versus r, λ or k respectively in the following figures.

We first consider the duration time as in Experiment 1. Theoretical arguments suggest that mean duration time should be a convex function of grain radius r, tending to infinity as $r \to 0$ and to a positive constant as $r \to \infty$. Simulations support this conclusion, but also show that the significant features of dependence occur over only a relatively small range of r.

In Figure 10 we plot the log(sample mean) over 200 realizations and the 95% confidence limits about the mean over the range $r \in [0.01, 0.5]$, from which we see that there are two peaks corresponding to two values of $r = 0.1$ and $r = 0.145$. This corresponds to the following observations: as r increases through $r = 0.1$ it becomes possible for one grain to cover at most 2 conditioning points, while as r increases through $r = 0.145$ it becomes possible for one grain to cover at most 4 conditioning points.

Now consider the dependence of mean duration time on intensity λ, as in Experiment 2. We gain a similar theoretical picture, except that for large intensity the mean duration time should increase from a minimum (since the algorithm needs more time to deal with more and more incidents of Y_A and $\mathcal{Z}_A$ processes as λ increases). Figure 11 plots the log(sample mean) over 200 realizations and 95% confidence limits about the mean, and confirms this argument from theory.

The sharp increase in duration time as intensity is reduced is mirrored by our ex-

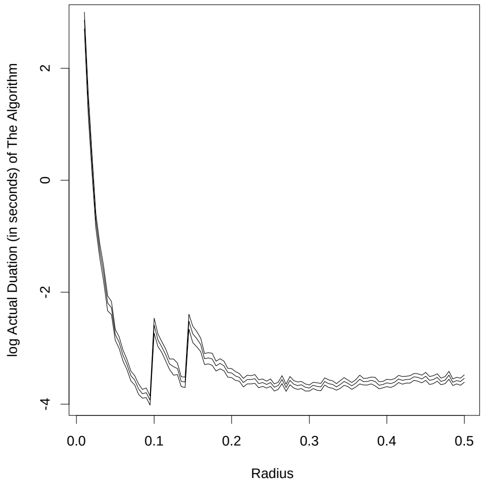

Figure 10: Graph of the log(mean actual duration) (in seconds) and 95% confidence limits over 200 realizations of the algorithm versus r for fixed value of $\lambda = 20$ and $k = 4$. The two peaks correspond to the two values of r when one grain can cover at most two and four conditioning points respectively.

perience of larger-scale simulations as displayed in Figures 1–6, particularly by the failure of the algorithm to produce coalescence for low intensity in such cases.

Things are much simpler in the case of dependence on number of conditioning points k as in Experiment 3. As in Experiment 3, fix $\lambda = 30, \ r = 0.04$. In this case, all the $\widetilde{Y}$ processes are independent with each other. Therefore, if we need h amount of time for coalescence for each conditioning point, then the total amount of time for coalescence for k conditioning points will be approximately kh. Figure 12 plots the log(sample mean) over 200 realizations and 95% confidence limits about the mean. This figure shows that the actual time increase as number of conditioning points increases, in agreement with the above.

7 Conclusion

In this paper we have presented a perfect simulation method for correlated Poisson random variables conditioned to be positive. Working in a general context (as opposed to the geometric context of [11]) has made the algorithm simpler to implement and to present. Moreover, although the method in this paper does not fall in the category of dominated CFTP defined in Kendall & Möller [10] nevertheless it is a modest variation of it. In this paper, "domination" does not occur with respect to the partial order defined in the state space, but the lower- and upper sandwich processes are "dominated" by a process in extended state space formed from Y_A and $\mathcal{Z}_A$ for all $A \subseteq S$.

While we cannot give theoretical bounds on the range of parameters involved in the particular kind of perfect simulation, it may be seen from the results of Sections

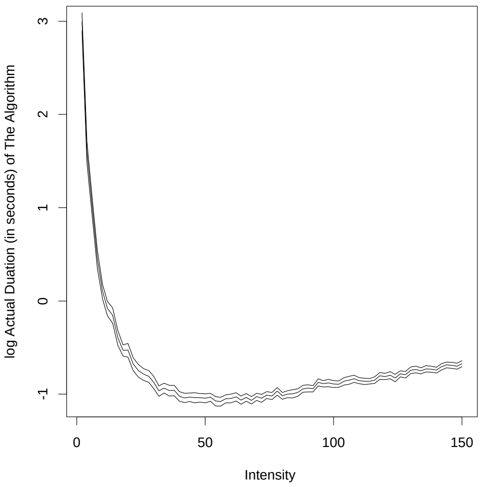

Figure 11: Graph of the log(mean actual duration) (in seconds) and 95% confidence limits over 200 realizations of the algorithm versus λ for fixed value of $r = 0.08$ and $k = 16$.

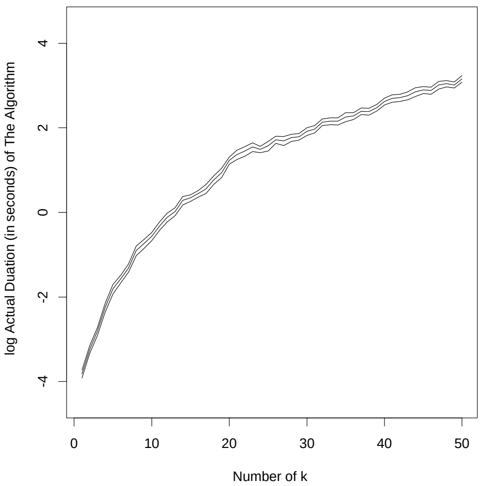

Figure 12: Graph of the log(mean actual duration) (in seconds) and 95% confidence limits over 200 realizations of the algorithm versus k for fixed value of $r = 0.04$ and $\lambda = 30$.

4.2 and 6 that low-intensity conditional Boolean models present challenges which this particular perfect simulation algorithm has not overcome. Of course there is a contrast here with conventional McMC methods: perfect simulation clearly indicates when it is failing to give a correct answer, while conventional McMC may fail without overt warning. This should not however be overstated. Firstly, there is a price to pay for CFTP [2], so that CFTP methods may fail when judiciously chosen McMC algorithms still give good answers. Secondly, CFTP exhibits bias if the user is liable to terminate runs early: this is overcome by Fill's rejection method [4] and this has been generalized to the point process context in [14] though, as Wilson has pointed out to one of us, a problem remains because of time-variability due to simulation of point processes as part of the point process algorithm. Despite these caveats, perfect simulation remains an attractive member of the statistician's tool-box, and the above algorithm delivers on its promise for medium and high intensity conditional Boolean models.

Further development of the particular problem of perfect simulation of conditional Boolean models leads to the consideration of conditioning constraints based on connectivity [12]. Perfect simulation in such situations is likely to prove very difficult, since the conditioning appears to involve arbitrarily large numbers of extra grains.

As with work on perfect simulation of Gibbs point processes [10, 9] we have taken care to test our simulation program carefully. This is a particular issue for perfect simulation implementations, which are inevitably more complicated than straightforward McMC algorithms, and raises interesting issues of how best to test correctness in such cases. In the above we have concentrated on straightforward χ^2 tests and compatibility with simple theoretical asymptotics: [9] gives details of a pseudo-likelihood approach which deals directly with the coupled random processes lying at the heart of the simulation. We hope to return to this issue in later work, as it has ramifications beyond perfect simulation in general issues of programming complicated simulations.

Acknowledgements

This research was supported by EPSRC grant GR/L56831 and by the EU TMR network ERB-FMRX-CT96-0095 on "Computational and Statistical methods for the analysis of spatial data".

References

[1] Barndorff-Nielsen, O.E. and Cox, D.R. (1989) *Asymptotic Techniques for use in Statistics*. Chapman and Hall.

[2] Burdzy, K. and Kendall, W.S. (1998) Efficient Markovian couplings: examples and counterexamples. *Research report 331, Department of Statistics, University of Warwick.* Submitted for publication.

[3] Cai, Y. and Kendall, W.S. (1999) Perfect implementation of simulation for conditioned Boolean model via correlated Poisson random variables. *Research report 350, Department of Statistics, University of Warwick.*

[4] Fill, J. (1998) An interruptible algorithm for exact sampling via Markov Chains. *Annals of Applied Probability*, 8:131–162.

[5] Grimmett, G.R. (1989) *Percolation*. Springer Verlag, New York.

[6] Kendall, W.S. (1997) On some weighted Boolean models. In D. Jeulin, editor, *Advances in Theory and Applications of Random Sets*, pages 105–120, Singapore. World Scientific Publishing Company.

[7] Kendall, W.S. (1997) Perfect simulation for spatial point processes. In *Bulletin ISI, 51st session proceedings, Istanbul (August 1997)*, volume 3, pages 163–166.

[8] Kendall, W.S. (1998) Perfect simulation for the area-interaction point process. In L. Accardi and C.C. Heyde, editors, *Probability Towards 2000*, pages 218–234, New York. Springer.

[9] Kendall, W.S. and Møller, J. (1999) Perfect implementation of a Metropolis-Hastings simulation of Markov point processes. *Research report 348, Department of Statistics, University of Warwick.*

[10] Kendall, W.S. and Møller, J. (1999) Perfect simulation of point processes using Metropolis-Hastings algorithms. *Research report 347, Department of Statistics, University of Warwick.*

[11] Kendall, W.S. and Thönnes, E. (1999) Perfect simulation in stochastic geometry. *Pattern Recognition*, 32(9):1569–1586.

[12] Lantuéjoul, C. (1997) Conditional simulation of object-based models. In D. Jeulin, editor, *Advances in Theory and Applications of Random Sets*, pages 271–288, Singapore. World Scientific Press.

[13] Propp, J.G. and Wilson, D.B. (1996) Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, 9:223–252.

[14] Thönnes, E. (1999) Perfect simulation of some point processes for the impatient user. *Advances in Applied Probability*, 31(1):69–87.

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.
293: *(with A.J. Baddeley and M.N.M. van Lieshout)* Quermass-interaction processes.
295: On some weighted Boolean models.
296: A diffusion model for Bookstein triangle shape.
301: COMPUTER ALGEBRA: an encyclopaedia article.
308: Perfect Simulation for Spatial Point Processes.
319: Geometry, statistics, and shape.

321: From Stochastic Parallel Transport to Harmonic Maps.
323: *(with E. Thönnes)* Perfect Simulation in Stochastic Geometry.
325: *(with J.M. Corcuera)* Riemannian barycentres and geodesic convexity.
327: Symbolic Itô calculus in *AXIOM*: an ongoing story.
328: *Itovsn3* in *AXIOM*: modules, algebras and stochastic differentials.
331: *(with K. Burdzy)* Efficient Markovian couplings: examples and counterexamples.
333: Stochastic calculus in *Mathematica*: software and examples.
341: Stationary countable dense random sets.
347: *(with J. Møller)* Perfect Metropolis-Hastings simulation of locally stable point processes.
348: *(with J. Møller)* Perfect implementation of a Metropolis-Hastings simulation of Markov point processes
349: *(with Y. Cai)* Perfect simulation for correlated Poisson random variables conditioned to be positive.
350: *(with Y. Cai)* Perfect implementation of simulation for conditioned Boolean Model via correlated Poisson random variables.

Also see the following related preprints

317: E. Thönnes: Perfect Simulation of some point processes for the impatient user.
334: M.N.M. van Lieshout and E. Thönnes: A Comparative Study on the Power of van Lieshout and Baddeley's J-function.

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.