Perfect simulation for interacting point processes, loss networks and Ising models

Pablo A. Ferrari[a,*], Roberto Fernández[b], Nancy L. Garcia[c]

[a]*Universidade de São Paulo, Brazil*
[b]*UPRES-A, CNRS 6085, Mathématiques, Site Colbert, Université de Rouen, F 76821 Mont Saint Aignan, Cedex, France*
[c]*IMECC, UNICAMP, Caixa Postal 6065, 13081-970 Campinas SP, Brazil*

Abstract

We present a perfect simulation algorithm for measures that are absolutely continuous with respect to some Poisson process and can be obtained as invariant measures of birth-and-death processes. Examples include area- and perimeter-interacting point processes (with stochastic grains), invariant measures of loss networks, and the Ising contour and random cluster models. The algorithm does not involve couplings of the process with different initial conditions and it is not tied up to monotonicity requirements. Furthermore, it directly provides perfect samples of finite windows of the *infinite-volume* measure, subjected to time and space "user-impatience bias". The algorithm is based on a two-step procedure: (i) a perfect-simulation scheme for a (finite and random) relevant portion of a (space–time) marked Poisson processes (free birth-and-death process, free loss networks), and (ii) a "cleaning" algorithm that trims out this process according to the interaction rules of the target process. The first step involves the perfect generation of "ancestors" of a given object, that is of predecessors that may have an influence on the birthrate under the target process. The second step, and hence the whole procedure, is feasible if these "ancestors" form a finite set with probability one. We present a sufficiency criteria for this condition, based on the absence of infinite clusters for an associated (backwards) oriented percolation model. The criteria is expressed in terms of the subcriticality of a majorizing multitype branching process, whose corresponding parameter yields bounds for errors due to space–time "user-impatience bias". The approach has previously been used, as an alternative to cluster

---

*Corresponding author. IME USP, Caixa Postal 66281, 05311-970 Sao Paulo, Brazil.
E-mail addresses: pablo@ime.usp.br (P.A. Ferrari), roberto.fernandez@univ-rouen.fr (R. Fernández), nancy@ime.unicamp.br (N.L. Garcia).
URLs: http://www.ime.usp.br/~pablo, http://www.univ-rouen.fr/upresa6085/Persopage/Fernandez/, http://www.ime.unicamp.br/~nancy

expansion techniques, to extract properties of the invariant measures involved. © 2002 Elsevier Science B.V. All rights reserved.

1. Introduction

Perfect simulations or *exact sampling* are labels for a recently developed set of techniques designed to produce output whose distribution is guaranteed to follow a given probability law. These techniques are particularly useful in relation with Markov Chain Monte Carlo, and their range of applicability is rapidly growing (see http://dimacs. rutgers.edu/~dbwilson/exact).

There are several techniques for perfect simulation of Markov processes. The most popular ones can be classified in two categories: Propp and Wilson's *Coupling from the Past* (CFTP) and Fill's Interruptible Algorithm. The first type of technique applies, in its original version, to invariant measures of Markov processes with a *finite coalescence time*. That is, of processes for which there exists a coupling among trajectories such that with probability one the trajectories starting from all possible initial states coalesce in a finite time. This includes all irreducible Markov processes with a finite state space. The coalescence property becomes difficult to check if the state space is very large. The problem can be overcome for processes with the following *monotonicity property*: there must exist a "maximal" and a "minimal" state and a coupling such that the coalescence of coupled trajectories starting from these two states implies the coalescence of all other trajectories ("monotone coupling"). Examples of processes with this property include Glauber dynamics of spin systems with the FKG property (Propp and Wilson, 1996). Other perfect simulation techniques are based on backward coupling of embedded regeneration times (Corcoran and Tweedie, 2001), finitary coding (Häggström and Steif, 2000; van den Berg and Steif, 1999), tempering algorithms, cluster representation of a Markov chain (Cai, 1999 preprint), regenerative construction (Comets et al., 2002), embedding the problem into the coloring of a graph (Fill and Huber, 2000), representation as an infinite mixture (Hobert and Robert, 2000) among others.

The basic CFTP algorithm, sometimes called *vertical CFTP*, is in general not applicable to processes with infinite state space. Indeed, most of them lack uniform ergodicity, a property shown (Foss and Tweedie, 1998) to be equivalent to the existence of a coalescence scheme as above. To cope with this situation, Kendall (1997, 1998) introduced *dominated CFTP* (also called *horizontal CFTP* and *coupling into and from the past*). This extension also requires the state space to have a partial order, as well as the existence of a monotone coupling among the target process and two reversible *sandwiching processes*, which must be easy to sample. Algorithms of this type are available for attractive point processes and, through a minor modification,

also for repulsive point processes (Kendall, 1998). Similarly, Häggström et al. (1999) combined ideas from CFTP and the two-component Gibbs sampler to perfect simulate from process in infinite spaces which do not have maximal (or minimal) elements.

The interruptible algorithm proposed by Fill (1998) (see also Thönnes, 1999) is an acceptance-rejection scheme which applies to invariant measures of Markov processes whose *time-reversed* process has a monotonicity property. Thus, its range of applicability overlaps with that of the CFTP algorithm at reversible monotone processes like Glauber dynamics of attractive automata or ferromagnetic spin systems and attractive point processes. Later developments have made Fill's algorithm applicable to other processes as well (Fill et al., 2000). An important advantage of this algorithm is that it is free of the so-called *impatient-user bias*: no bias is introduced if the user aborts a long run of the algorithm.

Kendall (1997, 1998) and Kendall and Møller (2000) proposed dominated CFTP schemes applicable to finite-volume measures which are absolutely continuous with respect to a finite Poisson point process and that can be obtained as the invariant measure of an interacting spatial birth-and-death process. These algorithms are based on two ingredients: (i) the "thinning" of a space–time marked Poisson process, and (ii) the coupled construction of upper and lower processes whose coalescence signals the output of a perfect sample. See the recent review of Møller (2000) for more references.

In this paper we propose a new perfect-simulation algorithm which applies to the same type of measures but has the following distinctive features:

- We sample directly from a time stationary realization of the process. There is no coalescence criterion, either between coupled realizations or between sandwiching processes. The scheme neither requires nor takes advantage of monotonicity properties.
- The scheme directly samples a finite window of the equilibrium measure in *infinite volume*. In contrast, Kendall (1997, 1998) focus on finite windows with fixed boundary conditions, and the infinite-volume limit requires an additional process of "perfect simulation in space". Our construction is in the spirit of the algorithms proposed by van den Berg and Steif (1999) and by Häggström and Steif (2000) to simulate infinite-volume measures for nearest-neighbor interactions in a lattice at high temperature or "high noise". In a sense, our algorithm is complementary to those, because it applies to regimes where they break down (e.g. at low temperature). We point out that before the arrival of the perfect simulation wave, Ferrari (1990), van den Berg (1993) and van den Berg and Maes (1994) have also proposed construction schemes for (infinite-volume) Gibbs measures of spin systems that can be easily transcribed into perfect-simulation algorithms.
- The construction has the added value of being a proven theoretical tool for the analysis of properties of the target measure. For instance, in Fernández et al. (1998, 2001) we used it to obtain mixing properties, finite-volume corrections and the asymptotic (in temperature) distribution of "defects" of the low-temperature Ising translation-invariant extremal measures. Ferrari and Garcia (1998) used a similar construction to show ergodicity of family of loss networks in *R*.

- More generally, the construction can be used as an alternative to the cluster-expansion technology (Brydges, 1986; Koteckỳ and Preiss, 1986; Dobrushin, 1996) for the study of spin models, at least those with a flipping symmetry. In fact, it seems to have a region of validity more extended than usual cluster-expansion approaches.

Our algorithm does involve the "thinning" of a marked Poisson process—the *free process*—which dominates the birth-and-death process, and it involves a time-backward and a time-forward sweep. But these procedures are performed in a form quite different from previous algorithms. The initial stage of our construction is done *towards the past*, starting with a finite window and retrospectively looking to *ancestors*, namely to those births in the past that could have (had) an influence on the current birth. The construction of the *clan of ancestors* constitutes the time-backward sweep of the algorithm. Once this clan is completely constructed, the algorithm proceeds in a time-forward fashion "cleaning up" successive generations according to appropriate penalization schemes. This "ancestors approach" offers some noteworthy advantages:

 (i) The algorithm constructs only the portion of the birth-and-death process *strictly* needed for the *final* window at $t = 0$. This economy has two important consequences: first, we can sample directly from the infinite-volume measure, without boundary effects. Second, the scheme works for point processes with quite general grain distribution, for instance chosen from an unbounded family of objects. As a consequence, it can be applied to the simulation of loss networks or of Peierls contours of the low-temperature Ising model.
 (ii) Perfect sampling is assured once the algorithm determines the "first" ancestors, that is those ancestors that themselves do not have ancestors. Thus, the algorithm determines by itself, in a single sweep, how far back into the past the simulation must go. This contrasts with usual CFTP schemes where algorithms may have to be iterated several times, going further and further into the past, until coalescence is achieved.

The relation "being ancestor of" induces a backwards in time *contact/oriented percolation* process. The algorithm is applicable as long as this oriented percolation process is subcritical. This implies the following limitations of our scheme:

- It works at *low density* of objects, at least in infinite volume. It may work at higher densities in finite volume, but we have not pursued this investigation.
- The birthrate of objects must be uniformly bounded. This is necessary to guarantee the existence of the dominating free birth-and-death process.

Our algorithm does not rely on any type of monotonicity. Therefore, for monotone systems our algorithm probably needs to go further back into the past than dominated CFTP or other schemes that exploit monotonicity. This loss in efficiency could be compensated, at least partially, by the "economy" and "single-sweep" features mentioned in (i) and (ii) above. As Prof. Kendall pointed out (private communication), dominated CFTP constructions "are wasteful in that they simulate past grains without regard to whether or not they are in the relevant percolation cluster, but efficient in that they use some kind of monotonicity to detect whether or not one needs to investigate further

back into the past". As a counterpart, insensitivity to monotonicity amounts to generality and versatility. In particular, this versatility could be used to offset the limitation imposed by the low-density constraint.

As an illustrative analogy, let us present a parallel with what happens in statistical mechanics, where studies usually rely on two types of methods: (I) those based on correlation inequalities, and (II) those expansion based. (These are not the only methods, others include exact solutions and more abstract arguments based on compactness or convexity.) Types (I) and (II) are mutually complementary. Correlation inequalities yield very strong results, often valid over whole regions in parameter space. Nevertheless, these results are rigid in that their validity depends on very precise symmetry (monotonicity) properties that can be easily destroyed by even infinitesimal perturbations. In contrast, expansion-based techniques are very versatile and robust. While their a priori range of convergence is limited—it is restricted to low densities or high temperatures—, it is often the case that suitable changes of variables place other regions of parameter space within the scope of expansion methods. For instance, expansion-based studies of high-temperature spin models work with (interacting) random walks (see, e.g. Dobrushin, 1996). Closer to, and above, the critical temperature alternative expansions are available (Olivieri and Picco, 1990; Fernández et al., 1992). At low temperature the right variables are the contours, and to get closer to (but below) the critical temperature coarse-grained contours are needed (Gawędzki et al., 1987). In fact, the belief is that there always are "good" variables that make everything diluted enough.

Existing dominated-CFTP algorithms are reminiscent of methods of type (I). They are very effective and apply for large intervals of rates. But they are also very specific, small alterations in the models could affect monotonicity and render an algorithm inapplicable. Our scheme could, perhaps, play a role similar to the studies of type (II) for simulation purposes. In this sense, it is crucial that monotonicity requirements be absent from the procedure. The goal is to change variables so to fall into a low-rate Poissonian (birth-and-death) process. Such changes will in general destroy any (obvious) monotonicity property. As an example, our scheme is capable to deal with Peierls contours, and hence to provide an exact-sampling algorithm for the low-temperature Ising model (of course, it is "exact" modulo time and space user-impatience). This is a region inaccessible to pre-existing algorithms. The use of other random objects (see the end of the previous paragraph) could yield analogous algorithms for other regions of the phase diagram.

The comparison of our algorithm with expansion methods is, in fact, more than just an analogy. Its theoretical basis has been used to construct an alternative to usual expansion methods in statistical mechanical (Fernández et al., 1998, 2001). This alternative has a provable region of validity that exceeds that of usual cluster-expansion treatments.

For the sake of completeness we start with the definitions of the most conspicuous space processes whose distributions we can perfect-simulate (Section 2). Examples include area- and perimeter-interacting point process, invariant measures of loss networks, the random cluster model and the contour representation of the "+" or "−" Ising measures at low temperature. Its relation with birth-and-death processes is discussed

immediately after (Section 3), together with the basic simulation approach for the latter. The perfect-simulation scheme is finally presented in Section 5. Its central piece is the time-backward construction of the clan of ancestors of a Poissonian birth-and-death process.

2. Point processes

Let **G** be a measurable space and v a Radon measure on **G**. Typically **G** is $\mathbb{R}^d$, $\mathbb{Z}^d$, $\mathbb{R}^d \times \mathbf{G}'$ or $\mathbb{Z}^d \times \mathbf{G}'$, where $\mathbf{G}'$ is a set of "animals" or "marks". Let

$$\mathscr{S} = \{\xi \in \mathbb{N}^{\mathbf{G}} : \xi(\gamma) > 0 \text{ only for a countable set of } \gamma \in \mathbf{G}\}.$$

A point process is a random element $N \in \mathscr{S}$. We denote with μ the law of a point process N. N is interpreted either as a random configuration of points or a random counting measure on **G**.

Poisson process. The first example is a Poisson process on **G** with intensity measure v. Its law is characterized by

$$\mu^0(N : N(B) = k) = \mathrm{e}^{-v(B)} v(B)^k / k!$$

for measurable $B \subset \mathbf{G}$; besides, under μ^0 $N(B_i)$ are independent if B_i are disjoint. When $\mathbf{G} = \mathbb{R}^d$ and $\Lambda \subset \mathbb{R}^d$ we call μ_Λ^0 the law of $N^0 \cap \Lambda$. We call a Poisson process on $\mathbf{G} = \mathbb{R}^d$ homogeneous when $v(\Lambda)$ is a function of $\ell(\Lambda)$, the Lebesgue measure of Λ. Similarly, when $\mathbf{G} = \mathbb{Z}^d$, the process is called homogeneous when $v(\Lambda)$ is a function of $|\Lambda|$, the number of points in $\mathbb{Z}^d \cap \Lambda$. In this case, the intensity is proportional to the Lebesgue (respectively, counting) measure and the factor of proportionality is called the *rate* which equals $v(\Lambda)$ for any Λ with unit Lebesgue measure (resp. counting measure).

Finite total rate. For future purposes we consider the case $v(\mathbb{R}^d \times \mathbb{R}^+) < \infty$; we interpret the last coordinate as time. One can compute the distribution of the (not necessarily finite) time τ_1, the smallest time-coordinate of the points (if any) of the process. Indeed, calling N the point Poisson process with rate v, for $0 \leqslant t \leqslant \infty$,

$$\mathbb{P}(\tau_1 > t) = \mathbb{P}(N(\mathbb{R}^d \times [0, t)) = 0) = \exp(-v(\mathbb{R}^d \times [0, t))). \qquad (2.1)$$

In the case of one-dimensional processes ($d = 0$) the above reads

$$\mathbb{P}(\tau_1 > t) = \mathbb{P}(N([0, t)) = 0) = \exp(-v[0, t)). \qquad (2.2)$$

In this paper we consider only point processes that are absolutely continuous with respect to a Poisson process with law μ^0. The law of these processes is characterized by

$$\mu(\mathrm{d}N) = \Psi(N) \mu^0(\mathrm{d}N),$$

where Ψ is the Radon-Nikodim derivative of μ with respect to μ^0.

A Poisson process that appears in the literature is the *germ-grain* Poisson process. In this case $\mathbf{G} = \mathbb{R}^d \times \mathscr{B}^0(\mathbb{R}^d)$, where $\mathscr{B}^0(\mathbb{R}^d)$ is the set of compact Borel sets of

$\mathbb{R}^d$. For each alive germ $x \in \mathbb{R}^d$, $g \in \mathscr{B}^0(\mathbb{R}^d)$ is the associated grain. Assume that the grains are determined by a random variable independent of the rest, given by a certain probability distribution π_x, which may depend on the germ location x. The intensity v is defined by

$$v(\mathrm{d}(x,g)) = f(x)\pi_x(\mathrm{d}g)\,\mathrm{d}x, \tag{2.3}$$

where $f(x)$ is the intensity of germs.

Area-interaction point processes. These processes have been introduced by Baddeley and van Lieshout (1995). This is a germ-grain process as defined above, but the grain shape is fixed and equal to a compact convex $G \subset \mathbb{R}^d$. We only need to keep track of the germs, so $\mathbf{G} = \mathbb{R}^d$. The intensity v is defined by $v(\mathrm{d}x) = \kappa\,\mathrm{d}x$, κ is a positive real number. The intersections of the grains determine a weight that corrects the otherwise Poissonian distribution of germs. The process is absolutely continuous with respect to the Poisson process μ^0 with intensity v. The law of the area-interaction process is defined for bounded windows $\Lambda \subset \mathbb{R}^d$ by

$$\mu_\Lambda(\mathrm{d}N) = \frac{\phi^{-m_d(N \oplus G)}}{Z_\Lambda(\kappa,\phi)}\,\mu_\Lambda^0(\mathrm{d}N), \tag{2.4}$$

where μ_Λ^0 is the law of the unit Poisson process in the box Λ, ϕ is a positive parameter, $Z_\Lambda(\phi)$ is a normalizing constant and $N \oplus G$ is the *coverage process* given by

$$N \oplus G := \bigcup_{x \in N} \{x + G\}. \tag{2.5}$$

Strauss process. The setup is the same as the area interaction process, but now the unit Poisson process is weighted according to an exponential of the number of pairs of points closer than a fixed threshold r. The measure is defined by

$$\mu_\Lambda(\mathrm{d}N) = \frac{1}{Z_\Lambda}\,\mathrm{e}^{\beta_1 N(\Lambda) + \beta_2 S(N,\Lambda)}\,\mu_\Lambda^0(\mathrm{d}N), \tag{2.6}$$

where $S(N,\Lambda)$ is the number of unordered pairs such that $\|x_i - x_j\| < r$. The case $\beta_2 > 0$ was introduced by Strauss (1975) to model the clustering of Californian red wood seedling around older stumps. However, (2.6) is not integrable in that case (see Kelly and Ripley, 1976).

Low-temperature Ising model. The well-known *Peierls contours* allow to map the "+" or "−" measures of the ferromagnetic Ising model at low temperature into an ensemble of objects—the contours—interacting only by perimeter exclusion. See, for instance, Section 5B of Dobrushin, 1996, for a concise and rigorous account of this mapping. The (discrete) set $\mathbf{G}$ consists of contours; these are hypersurfaces formed by a finite number of $(d-1)$-dimensional unit cubes—*links* for $d=2$, *plaquettes* for higher dimensions—centered at points of $\mathbb{Z}^d$ and perpendicular to the edges of the dual lattice $\mathbb{Z}^d + (\frac{1}{2},\ldots,\frac{1}{2})$. To each contour one can assign an "origin" in $\mathbb{Z}^d$ and say that two contours are equivalent if they coincide after a translation of the origin. Calling $\mathbf{G}'$ the set of contours modulus this class of equivalence, the set $\mathbf{G}$ can be expressed by $\mathbf{G} = \mathbb{Z}^d \times \mathbf{G}'$, where the first coordinate corresponds to the origin and

the second to the "shape" of the contour. Call two plaquettes *adjacent* if they share a $(d-2)$-dimensional face. A set of plaquettes, γ, is *connected* if for any two plaquettes in γ there exists a sequence of adjacent plaquettes in γ joining them. The set γ is *closed* if every $(d-2)$-dimensional face is covered by an even number of plaquettes in γ. *Contours* are connected and closed sets of plaquettes. For example, in two dimensions contours are closed polygonals. Two contours γ and θ are said to be *compatible* if no plaquette of γ is adjacent to a plaquette of θ. In two dimensions, therefore, contours are compatible if and only if they do not share the endpoint of a link. In three dimensions two compatible contours can share vertices, but not sides of plaquettes. Ising spin configurations in a bounded region with "+" (or "−") boundary condition are in one-to-one correspondence with families of pairwise compatible contours.

Let the *compatibility matrix* $I : \mathbf{G} \times \mathbf{G} \to \{0,1\}$ be defined by

$$I(\gamma,\theta) = \begin{cases} 0 & \text{if } \gamma \text{ and } \theta \text{ are compatible,} \\ 1 & \text{otherwise.} \end{cases} \tag{2.7}$$

The "Poisson process" μ_Λ^0 in $\mathbb{N}^{\mathbf{G}}$ is the product of Poisson measures whose γ-marginal is Poisson with mean

$$w(\gamma) := \exp(-\beta|\gamma|)$$

for $\gamma \subset \Lambda$. Here $|\gamma|$ stands for the number of plaquettes of γ. The intensity measure v is discrete: $\int_B \mathrm{d}v = v(B) = \sum_{\gamma \in B} w(\gamma)$.

Let μ_Λ be the measure defined by: for $\xi \in \{0,1\}^{\mathbf{G}}$, such that $\xi(\gamma) \leqslant 1\{\gamma \subset \Lambda\}$,

$$\mu_\Lambda(\xi) = \frac{1}{Z_\Lambda} \left(\prod_{\gamma,\theta:\xi(\gamma)\xi(\theta)=1} [1 - I(\gamma,\theta)] \right) \mu_\Lambda^0(\xi), \tag{2.8}$$

where β is a positive parameter called inverse temperature. The factor Z_Λ is just the normalization.

The random cluster model. Consider $\Lambda \subset \mathbb{Z}^d$ and let $\mathbb{B}(\Lambda) := \{(x,y) \in \Lambda \times \Lambda : |x-y|=1\}$ the set of bonds of Λ. A bond configuration $\zeta \in \{0,1\}^{\mathbb{B}(\Lambda)}$ is a function from $\mathbb{B}(\Lambda) \to \{0,1\}$. Bonds assigned 1 are called *open*, otherwise *closed*. A *cluster* of ζ is a set of sites connected with open bonds; sites surrounded only by closed bonds are clusters of size 1. Let $p \in [0,1]$ and $q > 0$ be parameters and define the finite volume measure

$$\varphi_\Lambda(\zeta) = \frac{1}{Z_\Lambda(p,q)} \, p^{O(\zeta)} (1-p)^{C(\zeta)} q^{L(\zeta)}, \tag{2.9}$$

where $O(\zeta)$ is the number of open bonds of ζ in Λ, $C(\zeta)$ is the number of closed bonds and $L(\zeta)$ is the number of clusters. The constant $Z_\Lambda(p,q)$ is the normalization. In other words, φ_Λ is absolutely continuous with respect to the product measure on $\mathbb{B}(\Lambda)$ with parameter p, with Radon-Nikodim derivative $q^{L(\zeta)}/Z_\Lambda(p,q)$. This model has been introduced by Fortuin and Kasteleyn (1972); a review can be found in Grimmett (1995).

Taking the connected sets of bonds as the basic objects, this model can be written as in (2.8). More precisely, if one says that two sets of bonds are *incompatible* whenever they share some vertices, and takes $\mathbf{G} = \{\gamma \subset \mathbb{B}(\Lambda): \gamma \text{ is finite and connected}\}$, the probability weights of the model can be written in the form (2.8) with the Poisson means

$$w(\gamma) = \left(\frac{p}{1-p}\right)^{B(\gamma)} \left(\frac{1}{q}\right)^{V(\gamma)-1}. \tag{2.10}$$

Here $B(\gamma)$ is the number of bonds of γ and $V(\gamma)$ the number of vertices in the extremes of the bonds of γ. [That is, $V(\gamma) = \#\{x \in \Lambda \text{ such that } (x,y) \in \gamma \text{ or } (y,x) \in \gamma \text{ for some } y \in \gamma\}$.] The transformation $Y: \{0,1\}^{\mathbb{B}(\Lambda)} \to \{0,1\}^{\mathbf{G}}$ defined by

$$Y(\zeta)(\gamma) = 1 \quad \Leftrightarrow \quad \gamma \text{ is a maximally connected set of open bonds of } \zeta$$

satisfies

$$\varphi_\Lambda(\zeta) = \mu_\Lambda(Y(\zeta)). \tag{2.11}$$

3. Birth-and-death processes

3.1. Definition and examples

The common feature linking all the spatial processes described in the previous section is that *all these distributions can be realized as invariant measures of spatial interacting birth-and-death processes.*

We consider the state space $\mathscr{S}$ of point configurations on $\mathbf{G}$ with a Radon measure v as in Section 2.

The *free* birth death process is characterized by the fact that individuals are born at intensity v and last for a random time exponentially distributed of mean one. The generator of the free process is

$$A^0 F(\eta) = \int_{\mathbf{G}} v(\mathrm{d}\gamma)[F(\eta + \delta_\gamma) - F(\eta)] + \sum_{\gamma \in \mathbf{G}: \eta(\gamma) > 0} \eta(\gamma)[F(\eta - \delta_\gamma) - F(\eta)]. \tag{3.1}$$

Here δ_γ is the configuration with only one point at γ and $(\eta + \xi)(\theta) = \eta(\theta) + \xi(\theta)$ (coordinatewise sum). The invariant (and reversible) measure for the free process is the Poisson process μ^0 with intensity v.

Let μ be a measure absolutely continuous with respect to μ^0 with Radon-Nikodim derivative Ψ: $\mathrm{d}\mu(\eta) = \Psi(\eta)\,\mathrm{d}\mu^0(\eta)$. Define

$$AF(\eta) = \int_{\mathbf{G}} v(\mathrm{d}\gamma) \frac{\Psi(\eta + \delta_\gamma)}{\Psi(\eta)} [F(\eta + \delta_\gamma) - F(\eta)]$$

$$+ \sum_{\gamma \in \mathbf{G}: \eta(\gamma) > 0} \eta(\gamma)[F(\eta - \delta_\gamma) - F(\eta)]. \tag{3.2}$$

The difference with the free process is that in the interacting process the rate of birth $v(\mathrm{d}\gamma)$ is corrected with the quotient $\Psi(\eta + \delta_\gamma)/\Psi(\eta)$, while the rate of death remains unaltered. The measure μ is reversible for the process with generator A. To better interpret this dynamics assume

$$\Delta_\Psi := \sup_{\eta,\gamma} \frac{\Psi(\eta + \delta_\gamma)}{\Psi(\eta)} < \infty$$

and define $M : \mathbf{G} \times \mathscr{S} \to [0,1]$ by

$$M(\gamma|\xi) = \frac{\Psi(\xi + \delta_\gamma)}{\Delta_\Psi \Psi(\xi)}. \tag{3.3}$$

If $\bar{v} = \Delta_\Psi v$, the generator can be rewritten as

$$AF(\eta) = \int_{\mathbf{G}} \bar{v}(\mathrm{d}\gamma) M(\gamma|\eta)[F(\eta + \delta_\gamma) - F(\eta)]$$

$$+ \sum_{\gamma \in \mathbf{G} : \eta(\gamma) > 0} \eta(\gamma)[F(\eta - \delta_\gamma) - F(\eta)]. \tag{3.4}$$

This dynamics has the following interpretation. When the current configuration of objects is ξ, object γ attempts to be born with rate $\bar{v}(\mathrm{d}\gamma)$ and is effectively born with probability $M(\gamma|\xi)$. The death rate of any object is one.

The interaction M induces naturally the notion of *incompatibility* between individuals. This is not a necessarily symmetric matrix $I : \mathbf{G} \times \mathbf{G} \to \{0,1\}$ defined by

$$I(\gamma, \theta) := 1 \left\{ \sup_\xi \{|M(\gamma|\xi) - M(\gamma|\xi + \delta_\theta)|\} > 0 \right\}, \tag{3.5}$$

where δ_θ is the configuration having unique individual θ and the supremum is taken over the set of those ξ such that ξ and $\xi + \delta_\theta$ are in the set of configurations (either $\{0,1\}^{\mathbf{G}}$ or $\mathbb{N}^{\mathbf{G}}$). The function $I(\gamma, \theta)$ indicates which individuals θ may have an influence in the birthrate of the individual γ. In the case of the spatial point processes described above the matrix is symmetric and given by $I((x,g),(x',g')) = 1\{(x + g) \cap (x' + g') \neq \emptyset\}$. If $I(\gamma, \theta) = 1$—that is, if the presence/absence of θ modifies the rate of birth of γ—we say that θ is *incompatible* with γ. For the Ising and random-cluster models one recovers (2.7).

Area interaction point processes. In the repulsive ($\phi < 1$) point process (2.4) we have $\Delta_\Psi = \kappa \phi^{-m_d(G)}$ and

$$\bar{v}(\mathrm{d}x) = \kappa \phi^{-m_d(G)} \mathrm{d}x, \tag{3.6}$$

$$M(x|\xi) = \phi^{m_d(G) - m_d((x+G) \backslash (\xi \oplus G))}. \tag{3.7}$$

For the attractive ($\phi > 1$) case,

$$\bar{v}(\mathrm{d}x) = \kappa \, \mathrm{d}x \tag{3.8}$$

$$M(x|\xi) = \phi^{-m_d((x+G) \backslash (\xi \oplus G))}. \tag{3.9}$$

Continuous unbounded one-dimensional loss network. A loss network models, for instance, the occurrence of calls in a communication network. Kelly (1991) reviews several discrete regimes and introduces the following continuous generalization. Callers are arranged along an infinitely long cable and each call between two points $s_1, s_2 \in \mathbb{R}$ on the cable involves just the segment between them. The cable has the capacity to carry simultaneously up to C calls past any point along its length. Hence, a call attempt between s_1 and $s_2 \in \mathbb{R}$, $s_1 < s_2$, is lost if past any point of the interval $[s_1, s_2]$ the cable is already carrying C calls. Calls are attempted with initial (leftmost) point following a space–time Poisson process with intensity $f(x)\,\mathrm{d}x$, and (space) lengths given by a distribution π, independent of its leftmost point, with finite mean ρ_1. The holding time of a call has exponential distribution with mean one. The location of a call, its length and its duration are independent.

In this case, the germs (x) are the leftmost points of calls and the grains (g) are segments with random lengths. This process can be viewed as a spatial birth and death process where we can take

$$\bar{v}(\mathrm{d}(x,g)) = f(x)\,\mathrm{d}x\pi(\mathrm{d}g), \tag{3.10}$$

and denoting $\xi(u) :=$ number of calls using point u,

$$M((x,g)|\xi) = 1\{(\xi + \delta_{\{x+g\}})(u) \leqslant C, \forall u \in \mathbb{R}\}. \tag{3.11}$$

Discrete processes. Free birth-and-death processes with a countable family of individuals **G** are simply the product of independent birth death processes labeled by each $\gamma \in \mathbf{G}$, with birthrates $w(\gamma)$ and deathrate equal to the number of alive individuals. Such a process exists without any requirement on the weights $w(\gamma)$; it is ergodic and its invariant distribution is the product of Poisson laws with mean $w(\gamma)$.

The (discrete) loss networks, the contour model and the animal version of the random cluster model of Section 2 are processes of this form where, in fact, the matrix M takes only two values, 0 and 1. That is, the interaction imposes a deterministic constraint. In particular, the interaction terms of the Ising-model, random cluster and loss networks have a simple product form

$$M(\gamma|\xi) = \prod_{\theta:\xi(\theta)\neq 0} [1 - I(\gamma,\theta)]. \tag{3.12}$$

Indeed, the Radon-Nikodim derivative is one for allowed configurations, hence the denominator in (3.3) is one.

3.2. Graphical construction

We proceed to the construction of the probability space where both the free process and the interacting birth death process will be constructed. Consider the countable family of random quartets $\{(\Gamma_i, T_i, S_i, Z_i): i \in \mathscr{I}\}$, with $\Gamma_i \in \mathbf{G}$, $T_i, S_i \in \mathbb{R}$, $Z_i \in [0,1]$ such that:

- The process $\{(\Gamma_i, T_i), i \in \mathscr{I}\}$ is a Poisson process on $\mathbf{G} \times \mathbb{R}$ with mean measure $v \times \ell$; ℓ is the Lebesgue measure in $\mathbb{R}$. This process determines the time and type of attempted births of individuals.

- S_i is exponentially distributed with mean 1. This variable will determine the lifetime of the ith attempted birth.
- Z_i is uniformly distributed in $(0, 1)$. This variable is called the *flag* or *mark* of the ith attempted birth and will be used together with the function $M(\cdot|\cdot)$ to decide if the attempted birth is actually a real birth.

Each triplet (Γ_i, T_i, S_i) can be visualized as a *cylinder* of (space) *basis* Γ_i, birth time T_i and lifetime S_i. The random set of marked cylinders is called

$$\mathbf{C} = \{(\Gamma_i \times [T_i, T_i + S_i] \ Z_i), i \in \mathscr{I}\}. \tag{3.13}$$

For a generic marked cylinder $C = (\Gamma \times [t, t+s], z) \in \mathbf{C}$, denote $\mathrm{birth}(C) = t$, $\mathrm{death}(C) = t + s$, $\mathrm{life}(C) = [t, t + s]$, $\mathrm{basis}(C) = \Gamma$ and $\mathrm{flag}(C) = z$.

The free process. The construction of $\mathbf{C}$ is time-translation invariant. Call

$$\xi_t(\gamma) := \#\{C \in \mathbf{C}: \ \mathrm{basis}(C) = \gamma; \ \mathrm{life}(C) \ni t\} \tag{3.14}$$

the set of individuals forming the sections of $\mathbf{C}$ at time t. All attempted births are actual births in this case. $\xi_t(\gamma)$ will be at most 1 in the continuous case, but could be bigger in the discrete case. $(\xi_t : t \in \mathbb{R})$ constitutes a stationary free birth death process with generator (3.1). The marginal law of ξ_t is μ^0, the Poisson process with intensity v.

Likewise, one can define the free process on $\mathbb{R}^d \times [0, \infty)$ with initial configuration of individuals

$$\xi_0 := \{\gamma_1^0, \gamma_2^0, \ldots\}. \tag{3.15}$$

For this associate cylinders to the initial configuration:

$$\mathbf{C}_0(\xi_0) := \{(\gamma_i^0 \times [0, 0 + S_i^0], Z_i^0), i \in \mathbb{Z}\}, \tag{3.16}$$

where S_i^0 and Z_i^0 are independent and independent-of-everything random variables whose distributions are, respectively, $\mathrm{Exp}(1)$ and $U(0, 1)$. Define the subset of cylinders born between 0 and t:

$$\mathbf{C}_{[0,t]} := \{C \in \mathbf{C}: \mathrm{birth}(C) \in [0, t]\}.$$

Then, the process defined at time t by

$$\xi_t(\gamma) := \#\{C \in \mathbf{C}_{[0,t]} \cup \mathbf{C}_0(\xi_0): \mathrm{basis}(C) = \gamma; \mathrm{life}(C) \ni t\} \tag{3.17}$$

has initial configuration ξ_0 and generator (3.1).

Interacting processes. The absolute continuity with respect to the free process, embodied in the generator (3.4), suggests a simple alteration to the previous construction to pass to an interacting birth-and-death process: the *attempted births* become actual births only if an additional (generally stochastic) test is passed. This test is determined by the factor M of the rate densities. The interacting process is, therefore, obtained as a "thinning" or "trimming" of the free process.

The formalization of this intuitively simple idea is easy for finite windows, but more delicate for the infinite-volume process. We discuss the former case first.

3.2.1. Finite-volume construction

To construct a birth-and-death process ξ_t with rate density $\bar{v}(d\gamma)M(\gamma|\xi)$, for individuals within a finite space-region Λ and for a finite time interval $[t_0, t_{\mathrm{fin}}]$, one proceeds as follows:

(1) Run the free process with rate density $\bar{v}$ starting from the initial cylinders $\mathbf{C}_0$. If $M(\gamma|\xi)$ is deterministic—for instance forbidding individuals to overlap—the initial configuration is assumed to satisfy the corresponding constraint.
(2) Each death happening before reaching an event of the free process causes the corresponding updating of ξ_t, by taking the corresponding individual out of ξ_t.
(3) When the free process yields a first event $(\gamma_1, t_1, s_1, z_1)$, this event is considered an attempted birth. To decide, one looks to the set ξ_{t_1-} of alive individuals (ξ_{t_1-} is equal to ξ^0 minus the initial individuals with lifetime smaller than t_1). If

$$z_1 < M(\gamma_1|\xi_{t_1-}) \tag{3.18}$$

the cylinder is allowed to be born and the individual γ_1 is included in the configuration ξ_{t_1}; otherwise it is ignored and ξ_{t_1} is set equal to ξ_{t_1-}.
(4) Now iterate the procedure, that is, repeat the previous two steps shifting subscripts $1 \to 2$ and $0 \to 1$. Continue in this way until reaching an attempted birth beyond t_{fin}.

3.2.2. Two-sweep finite-volume construction

The visualization in terms of cylinders suggests an alternative implementation as a two-sweep scheme: in the first sweep one generates free cylinders by running the free process from t_0 to t_{fin}, while in the second sweep a decision is made on which cylinders are kept and which are erased. The set of *kept cylinders* includes, by definition, all initial cylinders while successive additions must pass the test (3.18). We call $\mathbf{K}_{[0,t]}(\Lambda, \xi_0)$ the resultant set of kept cylinders in the construction of Section 3.2.1. The configuration of the process at time t with initial configuration ξ_0 is then given by the projection of the bases of the alive kept cylinders at that time:

$$\eta_t(\gamma) = \#\{C \in \mathbf{K}_{[0,t]}(\Lambda, \xi_0) \text{ with basis } \gamma \text{ and alive at } t\}. \tag{3.19}$$

3.2.3. Finite-volume time-stationary construction

The construction can be also performed in a stationary manner for $t \in \mathbb{R}$. Indeed, since in a finite window the number of alive individuals is finite (with probability one), there exist random times $\{\tau_j \in \mathbb{R}: j \in \mathbb{Z}\}$, such that (a) $\tau_j \to \pm\infty$ for $j \to \pm\infty$ and (b) $(\bigcup_i [T_i, T_i + S_i]) \cap (\bigcup_j \{\tau_j\}) = \emptyset$. In words, at each τ_j no cylinder is alive. The above selection of kept cylinders can then be performed independently in each of the random intervals $[\tau_i, \tau_{i+1})$. This stationary construction is particularly useful to study properties of the invariant measure μ_Λ. In fact, calling $\mathbf{K}(\Lambda)$ the (time stationary random) set of kept cylinders, the law of

$$\eta_t(\gamma) = \#\{C \in \mathbf{K}(\Lambda) \text{ with basis } \gamma \text{ and alive at } t\} \tag{3.20}$$

is exactly μ_Λ.

3.2.4. Infinite-volume construction

None of the finite-volume procedures discussed above can be directly implemented to construct the process in infinite volume. On the one hand, the scheme proposed for finite time intervals is not applicable to infinite volume because it is not possible to decide which is the first mark in time. On the other hand, the stationary construction is also not feasible because in infinite volume there are cylinders alive at all times. This last objection, however, may play no role if one only focuses on a family of cylinders intersecting a *finite* set and tries to decide which of them should be erased and which ones kept.

According to the previous discussion, to decide whether a cylinder $C \in \mathbf{C}$ is kept, one has to look at the set of cylinders C' (born before C and) alive at the birth time of C whose basis are incompatible with the basis of C in the sense of (3.5). Let us call this set the first generation of *ancestors* of C and denote it $\mathbf{A}_1^C$. Once we determine which of these ancestors are alive, the decision on whether to keep C or not requires only a single application of the test (3.18). However, to decide which of these ancestors are alive we have to work with the second generation of ancestors of C, that is, with the ancestors of the ancestors. Recursively, we find ourselves having to deal with all generations of ancestors of C. Let us call the union of all generations of ancestors of C the *clan of ancestors* of C, and denote it $\mathbf{A}^C = \bigcup_{n \geqslant 1} \mathbf{A}_n^C$, where $\mathbf{A}_n^C$ is the set of ancestors in the nth generation. These sets may contain cylinders in $\mathbf{C}_0(\xi_0)$. The procedure for deciding whether to keep or to erase C can be univocally defined if *the clan of ancestors of C is finite*.

This picture makes it apparent that an infinite-volume process given initial starting conditions (i.e. for a *finite time interval*) exists as long as there are no explosions, that is, as long as no cylinder can develop infinitely many ancestors in a finite time. Furthermore, there exists a unique stationary process (for *infinite time intervals*) if *all* clans of ancestors are finite with (free process) probability one.

Theorem 1. (i) *If with probability one $\mathbf{A}^C \cap \mathbf{C}_{[0,t]}$ is finite for every cylinder C alive at time t, for any $t > 0$, then the birth-and-death process with the generator* (3.2) *and initial condition η_0 is obtained by performing the two-sweep construction of Section* (3.2.2) *on each set $(\mathbf{A}^C \cap \mathbf{C}_{[0,t]}) \cup \mathbf{C}_0(\eta_0)$ and taking the projections*

$$\eta_t(\gamma) = \#\{C \in \mathbf{K}_{[0,t]}(\eta_0) \text{ with basis } \gamma \text{ and alive at } t\}, \qquad (3.21)$$

where $\mathbf{K}_{[0,t]}(\eta_0)$ is the resulting set of kept cylinders.

(ii) *If with probability one $\mathbf{A}^C$ is finite for every cylinder C, then the stationary birth-and-death process with the generator* (3.2) *can be constructed for $t \in \mathbb{R}$ by performing the two-sweep construction of Section* (3.2.2) *on each set $\mathbf{A}^C$ and taking the projections*

$$\eta_t(\gamma) = \#\{C \in \mathbf{K} \text{ with basis } \gamma \text{ and alive at } t\}. \qquad (3.22)$$

Moreover, the marginal distribution of η_t is the stationary measure μ.

A proof of this theorem is presented in Appendix A.

4. Oriented percolation and branching processes

To determine the conditions allowing the construction of Theorem 1, we point out that the relation "being ancestor of" gives rise to a model of *oriented percolation*. We call it *backwards oriented percolation* to emphasize the fact that it is defined by only looking into the past. The finite-time construction is possible if there is no cluster with infinitely many members in a finite-time slice, while the feasibility of the infinite-time construction requires the absence of a percolation cluster reaching to time $-\infty$. As usual in oriented percolation problems, it is useful to work with a majorizing *multitype branching process*. In this process the offspring distribution of a cylinder C has the same (marginal) law as the distribution of $\mathbf{A}_1^C$, but the branches behave independently. The problem is then reduced to determine conditions guaranteeing the finiteness of the clan of branching ancestors. Hence it is sufficient to show in the finite-time case that the branching process does not explode, while in the infinite-time case we need to prove that the branching process is subcritical. Factorization makes these tasks easier.

Let us give sufficient conditions on the dominating branching for the different processes listed in Section 2.

Discrete processes. If the family of individuals $\mathbf{G}$ is countable, the free birth-and-death process is the product over $\gamma \in \mathbf{G}$ of independent marked Poisson processes. The construction of the interacting processes is an obvious adaptation of the procedure of the continuous case. For the infinite-volume process one relies on the properties of the backwards oriented percolation model of cylinders defined by the oriented bonds $C \to C'$ if C' is an ancestor of C, that is if the basis of C and C' intersect and C' is alive when C is born. Let $m(\gamma, \theta)$ be the mean number of cylinders of basis θ in the first generation of a cylinder of basis γ. These are cylinders born at negative times $-t$ and have a lifetime at least t, so they survive to intersect the grain born at time zero. Its average number is, therefore,

$$m(\gamma, \theta) = w(\theta)I(\gamma, \theta) \left[\int_{-\infty}^{0} dt \int_{t}^{\infty} ds \, e^{-s} \right] = w(\theta)I(\gamma, \theta) \cdot 1. \tag{4.1}$$

Define $m^n(\gamma, \theta)$ as the mean number of cylinders of basis θ incompatible with a cylinder of basis γ in the nth generation of ancestors, m^n is the matrix product of m by itself n times. The condition for absence of oriented percolation is

$$\sum_{n \geqslant 1} \sum_{\theta} m^n(\gamma, \theta) < \infty \tag{4.2}$$

for all γ. For any function "size" $q : \mathbf{G} \to \mathbb{R}^+$, such that $\inf_\gamma q(\gamma) \geqslant 1$, as in Lemma 5.15 of Fernández et al. (2001), calling

$$\alpha_q := \sup_\gamma \frac{1}{q(\gamma)} \sum_\theta q(\theta)m(\gamma, \theta), \tag{4.3}$$

we have

$$\sum_\theta m^n(\gamma, \theta) \leqslant \alpha_q^n q(\gamma). \tag{4.4}$$

The form of measuring this "size" depends on the process in question, but usually there is an obvious prescription. For instance, for the loss networks, the Peierls contours and the random clusters model this measure is just the length of the call, the number of plaquettes of the contour or the number of points of the cluster. The (infinite-volume) birth-and-death process—and hence the corresponding loss network—exists for finite time intervals if $\alpha_q < \infty$, while $\alpha_q < 1$ is a sufficient condition to be an ergodic infinite-time process.

Area interaction-point process. Here $m(x,.)$ is a measure on $\mathbf{G} = \mathbb{R}^d$; $m(x, \mathrm{d}y)$ represents the rate at which cylinders with basis centered at y appear. Consider a germ x_0, which, by space–time-invariance, can be placed at the origin and assumed to be born at time zero. Its ancestors are all cylinders whose bases involve germs located in $\partial G := \{x : (x + G) \cap G \neq \emptyset\}$. Therefore, as in (4.1),

$$m(x, \mathbf{G}) = v(\partial G) \cdot 1. \tag{4.5}$$

We conclude that the corresponding birth-and-death process exists for finite times as long as $v(\partial G) < \infty$ and, if furthermore,

$$v(\partial G) < 1 \tag{4.6}$$

then there is an stationary ergodic process, absolutely continuous respect to the free process, having as invariant measure the corresponding point process of Section 2.

The argument also works if $G \in \mathbf{G}$ is a random set chosen independently of everything as in (2.3). Recall $f(x)$ is the intensity of germs and π_x is the distribution of the grain centered in x. Let

$$m(G, \mathrm{d}H) = \int \mathrm{d}x \, f(x) I(G, H) \, \pi_x(\mathrm{d}H) \tag{4.7}$$

the rate at which individuals H having an influence in the birthrate of G appear (see (3.5) for the definition of I). This implies that the mean number of individuals in the first generation of ancestors of G is $m(G, \mathbf{G})$. Let the "matrix product" m^n be defined inductively by $m^1 = m$ and

$$m^n(G, \mathrm{d}H) := \int_{\mathbf{G}} m^{n-1}(G, \mathrm{d}K) \, m(K, \mathrm{d}H) \quad n > 1. \tag{4.8}$$

As for $n = 1$, $m^n(G, \mathbf{G})$ is the mean number of individuals in the nth generation of ancestors of G. The ergodicity of the process is implied by

$$\sum_{n \geqslant 1} m^n(G, \mathbf{G}) < \infty \tag{4.9}$$

for all individual G. A sufficient condition for (4.9) is

$$\alpha_q := \sup_G \frac{1}{q(G)} \int_{\mathbb{R}^d} f(x) \, \mathrm{d}x \int_{\mathbf{G}} \pi_x(\mathrm{d}H) q(H) I(G, H) < 1 \tag{4.10}$$

for some function $q : \mathscr{B}(\mathbb{R}^d) \to \mathbb{R}^+$ satisfying $\inf_G q(G) \geqslant 1$. Indeed, it can be proven as in Lemma 5.15 in Fernández et al. (2001)

$$m^n(G, \mathbf{G}) \leqslant q(G) \alpha_q^n. \tag{4.11}$$

Strictly speaking the above statements have been rigorously proven only for the discretized version of the models.

Loss networks. The calls of the loss networks can be interpreted as germ-grains. For instance, in the one-dimensional case, the germs are the leftmost points of the calls and the grains are segments with random lengths.

A particular case where one can explicitly compute the sufficient condition is the one-dimensional continuous loss networks of Section 3.1. Assume, in general, that the leftmost points of calls appear with rate $f(x)$ and that call lengths are given by a distribution π independent of x. We only require the latter to have a finite mean ρ_1. Consider a germ sitting at the origin, that is a call stretching from the origin to the right, born at time zero. Its ancestors correspond to cylinders with sufficient lifetime and with bases given by either calls starting at negative sites and passing through the origin, or calls of arbitrary length originating within the sites occupied by the initial call. Therefore, the α_q in (4.10) for the case $q(L) \equiv 1$ turns to be

$$\alpha_q = \sup_L \left(\int_{-\infty}^0 \pi\{L > -x\} f(x)\, dx + \int_0^L dx\, f(x) \right). \tag{4.12}$$

In the homogeneous case ($f(x) \equiv \kappa$) this gives the following condition for ergodicity:

$$\kappa \left(\rho_1 + \sup_L L \right) < 1. \tag{4.13}$$

A simple computation shows that choosing $q(L) = \max(L, 1)$ gives $\alpha_q \leqslant \kappa(\rho_2 + \rho_1 + 1)$, where ρ_1 and ρ_2 are the first and second moment of the distribution π, respectively. This gives the following sufficient condition for ergodicity:

$$\kappa(\rho_2 + \rho_1 + 1) < 1. \tag{4.14}$$

Máric (2002) improved this bound to

$$\kappa(\sqrt{\rho_2} + \rho_1) < 1. \tag{4.15}$$

We remark that to obtain these conditions it was important to consider only *oriented* percolation. The analogous conditions obtained by considering unoriented percolation of cylinders are far more restrictive.

5. Perfect simulation of invariant measures of birth-and-death processes

The main issue of this section is a construction of the set $\mathbf{A}^{\Lambda,0}$ formed by the cylinders with bases intersecting the space–time set $\Lambda \times \{0\}$ ("cylinders alive at time 0") and their clans of ancestors. This is a problem of simulation of cylinders generated by the *free* process. Once these clans are perfectly simulated, it is only necessary to apply the *deterministic* "cleaning procedure", based on the test (3.18), to obtain a perfect sample of the interacting process. The scheme is feasible if these clans are finite with probability one, a fact valid under conditions like (4.6), (4.9) or $\alpha_q < 1$, where α_q is defined in (4.3) for the discrete case and (4.10) for the area interaction process.

We propose a non-homogeneous time-backwards construction of the clan based on a result proven in Section 4.5.1 of Fernández et al. (2001). It is shown there that the clan of ancestors of a family of cylinders can be obtained combing *back* in time and *generating births* of ancestors with an appropriate rate. Alternatively, one could use the fact that the law of **C** is time-reflexion invariant, to *generate deaths* of ancestors. This is simple in the area-interaction process with a fixed grain, but it is not Markovian and more involved in the infinite case. This approach was proposed by one of the referees and developed by Garcia (2000).

For concreteness, let us discuss our scheme for individuals living in $\mathbb{R}^d$ or $\mathbb{Z}^d$. The birthrate of a new cylinder to be added to the clan is equal to the rate density of the free process multiplied by an exponential time factor ensuring that the ancestor has a lifespan large enough to actually be an ancestor. This time factor involves the time–distance to the birth of existing cylinders, which can be expressed through the following function. For a finite region Λ and a finite set of cylinders **H**, let the set of bases of the potential ancestors of **H** and $\Lambda \times \{0\}$ be

$$\mathbf{G}(\mathbf{H}, \Lambda) := \{\theta \in \mathbf{G} \colon I(\mathrm{basis}(C'), \theta) = 1, \text{ for some } C' \in \mathbf{H}\}$$

$$\cup \{\theta \in \mathbf{G} \colon \theta \cap \Lambda \neq \emptyset\} \tag{5.1}$$

and for a given individual $\theta \in \mathbf{G}(\mathbf{H}, \Lambda)$,

$$\mathrm{TI}(\mathbf{H}, \Lambda, \theta) = \min\{\mathrm{birth}(C') \colon C' \in \mathbf{H}, \ I(\mathrm{basis}(C'), \theta) = 1\} \tag{5.2}$$

with the convention $\min \emptyset = 0$. By definition, $\mathrm{TI}(\mathbf{H}, \Lambda, \theta) \leqslant 0$.

Theorem 2. *The clan $\mathbf{A}^{\Lambda, 0}$ is the limit as $t \to \infty$ of a process $\mathbf{A}_t$, defined by the initial condition $\mathbf{A}_0 = \emptyset$ and the evolution equation*

$$\mathbb{E}\left(\frac{\mathrm{d}F(\mathbf{A}_t)}{\mathrm{d}t} \,\middle|\, \mathbf{A}_s, \ 0 \leqslant s \leqslant t \right)$$

$$= \int_{\mathbf{G}(\mathbf{A}_t, \Lambda)} v(\mathrm{d}\theta) \int_{t+\mathrm{TI}(\mathbf{A}_t, \Lambda, \theta)}^{\infty} \mathrm{d}s \, \mathrm{e}^{-s} [F(\mathbf{A}_t \cup (\theta, -t, s)) - F(\mathbf{A}_t)]. \tag{5.3}$$

Here F is an arbitrary function depending on a finite number of individuals intersecting Λ and we have denoted $(\theta, -t, s)$ the cylinder of base θ, born at time $-t$ and with lifetime s.

For completeness, a proof of this theorem is presented in Appendix A. For the free discrete loss-network (contours, random cluster) processes,

$$\int_{\mathbf{G}(\mathbf{A}_t, \Lambda)} v(\mathrm{d}\theta) F(\theta) = \sum_{\theta \in \mathbf{G}(\mathbf{A}_t, \Lambda)} w(\theta) F(\theta) \tag{5.4}$$

while for free birth-and-death processes on $\mathbb{R}^d$,

$$\int_{\mathbf{G}(\mathbf{A}_t, \Lambda)} v(\mathrm{d}\theta) F(\theta) = \int_{\mathbb{R}^d} f(x) \, \mathrm{d}x \int_{\mathbf{G}_x} \pi_x(\mathrm{d}g) 1\{x \oplus g \in \mathbf{G}(\mathbf{A}_t, \Lambda)\} F(x, g).$$

Notice that $\mathbf{A}_t$ is a monotone process $(\mathbf{A}_t \subset \mathbf{A}_{t+s})$ in which at time t only cylinders in $\mathbf{G}(\mathbf{A}_t, \Lambda)$ can be included. The inclusion of a cylinder born at time $-t$ requires that either (a) its basis is incompatible with that of some cylinder born later and its lifespan reaches the birth time of such cylinder, or (b) its basis is compatible with those of all cylinders born later, but it intersects Λ and the cylinder survives up to time equal zero. The last condition is ensured via the convention $\min \emptyset = 0$ in the definition of TI.

Algorithm to construct the backwards clan of a finite region. The combination of (5.1)/(5.3) can be translated into the following explicit algorithm. We do it first for the case of countable number of individuals and indicate at the end of this section how to proceed in the continuous case. To generate $\mathbf{A}^{\Lambda,0}$:

(1) Start with $\tau = 0$ and $\mathbf{H} = \emptyset$.
(2) Let $\mathbf{H}$ be the current set of cylinders and τ the current $-\min\{\text{birth}(C)\colon C \in \mathbf{H}\}$. For each $\gamma \in \mathbf{G}(\mathbf{H}, \Lambda)$ generate an independent realization of the first time $\tau_1(\gamma)$ of the non-homogeneous Poisson process in $\mathbb{R}$ with intensity

$$\lambda_\gamma(\mathrm{d}s) := w(\gamma)\mathrm{e}^{-s+\mathrm{TI}(\mathbf{H},\Lambda,\gamma)}1\{s > \tau\}\,\mathrm{d}s. \tag{5.5}$$

Notice that $\tau_1(\gamma)$ may be infinity.
(3) Order the set $\{\tau_1(\gamma)\colon \gamma \in \mathbf{G}(\mathbf{H}, \Lambda)\}$. Let $\tilde{\tau}$ be the infimum of this set. This is well defined and strictly positive because the condition $\alpha < 1$ guarantees that the total rate $\sum_{\gamma \in \mathbf{G}(\mathbf{H},\Lambda)} \int_{\mathbb{R}^+} \lambda_\gamma(\mathrm{d}s) < \infty$.
(4) If $\tilde{\tau} < \infty$, call γ_1 the basis corresponding to the minimum (i.e. $\tau_1(\gamma_1) = \tilde{\tau}$). Update $\tau \leftarrow \tilde{\tau}$ and $\mathbf{H} \leftarrow \mathbf{H} \cup \{(\gamma_1, -\tau, \tau + \mathrm{TI}(\mathbf{H}, \Lambda, \gamma) + R_1)\}$, where R_1 is an exponential random variable with rate 1 independent of everything. In the sequel ignore the set $\{\tau_1(\gamma)\colon I(\gamma, \gamma_1) = 1\}$ (we can reuse the remaining τ_1) and go to (2).
(5) If $\tilde{\tau} = \infty$ set $\mathbf{A}^{\Lambda,0} = \mathbf{H}$ and stop. By Theorem 2 the distribution of the set $\mathbf{A}^{\Lambda,0}$ so generated is exactly that of the free birth-and-death process.

If τ_i are the successive times of jump of $\mathbf{A}_t$, then $\mathbf{H}_i = \mathbf{A}_{\tau_i}$ have the same distribution as the ith iterate of the above algorithm.

In the continuous case, time and space cannot be in general separated. Instead of steps (2) and (3) above we must consider the events (γ, s) of a Poisson process on $\mathbf{G} \times \mathbb{R}^+$ with intensity

$$\lambda(\mathrm{d}(\gamma, s)) = v(\mathrm{d}\gamma)\,\mathrm{e}^{-s+\mathrm{TI}(\mathbf{H},\Lambda,\gamma)}1\{s > \tau\}1\{\gamma \in \mathbf{G}(\mathbf{H}, \Lambda)\}\,\mathrm{d}s. \tag{5.6}$$

For a finite window Λ the total rate is finite, hence these events can be well ordered by looking to the time coordinate. If the set of these events is not empty, we take $\tilde{\tau}$ to be the minimal time coordinate (it is strictly positive with probability one) and denote γ_1 the associated individual. If the Poisson process with rate density (5.6) yields no event we take $\tilde{\tau} = \infty$. We then continue as in (4).

This algorithm plus the subsequent "cleaning algorithm" constitutes our perfect simulation scheme.

The cleaning algorithm. Let $\mathbf{A}^{\Lambda,0}$ be the clan of the cylinders whose life contains time 0 and basis intersects Λ. The following algorithm shows how to construct inductively the set $\mathbf{K}^{\Lambda,0}$ of kept cylinders.

(1) Start with $\mathbf{H} = \mathbf{A}^{\Lambda,0}$ and $\mathbf{K} = \emptyset$ ($\mathbf{H}$ is formed by the cylinders to be tested and $\mathbf{K}$ by those already kept).

(2) If $\mathbf{H}$ is empty go to 5. If not, order the cylinders of $\mathbf{H}$ by time of birth. Let C_1 be the first of those cylinders; call γ_1 its basis and τ_1 its birth time. Let ξ_1 be the set of bases of the cylinders in $\mathbf{K}$ alive at τ_1 which are incompatible with the basis of C_1. Let Z_1 be a random variable uniformly distributed in $[0,1]$ independent of everything.

(3) If $Z_1 < M(\gamma_1|\xi_1)$, then update: $\mathbf{H} \leftarrow \mathbf{H} \setminus \{C_1\}$, $\mathbf{K} \leftarrow \mathbf{K} \cup \{C_1\}$. Go to 2.

(4) If $Z_1 > M(\gamma_1|\xi_1)$, then update: $\mathbf{H} \leftarrow \mathbf{H} \setminus \{C_1\}$. Go to 2.

(5) Set $\mathbf{K}^{\Lambda,0} = \mathbf{K}$ and stop. By Theorem 1(ii) the distribution of this clan $\mathbf{K}^{\Lambda,0}$ is exactly that of the interacting birth-and-death process.

Algorithm to simulate a finite window of μ. This is the easiest part. Once the set $\mathbf{K}^{\Lambda,0}$ of kept cylinders has been determined, take the configuration η defined by

$$\eta(\gamma) = \sum_{C \in \mathbf{K}^{\Lambda,0}} 1\{C \text{ has basis } \gamma \text{ and life containing } 0\} \tag{5.7}$$

for γ intersecting Λ. This configuration has the marginal distribution of the infinite-volume measure μ on the (not necessarily finite) set $\mathbf{G}_\Lambda = \mathbb{N}^{\{\theta \in \mathbf{G}: \theta \cap \Lambda \neq \emptyset\}}$. This fact is guaranteed by Theorem 1(ii).

6. Errors in perfect simulation?

Even in finite volume, perfect-simulation algorithms are subjected to error. In general terms, a perfect-simulation algorithm of a measure μ on a set $\mathscr{X}$ is a function $\Phi : [0,1]^{\mathbb{N}} \to \mathscr{X}$, such that, if $(U_n)_{n \in \mathbb{N}}$ is a sequence of i.i.d. uniform in $[0,1]$ random variables, there exists a stopping time T for (U_n) such that Φ depends only on the first T coordinates of $(U_1, U_2, \ldots)$ and

$$\mathbb{P}(\Phi(U_1, \ldots, U_T) \in A) = \mu(A). \tag{6.1}$$

The CFTP algorithm, for instance, stops when a random value t is found such that the different copies of the algorithm coupled from time $-t$ started with all possible initial conditions attain the same configuration at time 0. Finding t requires the use of a random number $T(t)$ of uniform random variables, which must be less than $S =$ "the maximum time left in order to have the results ready for the next congress", for instance. Thus, one actually samples from the distribution defined by

$$\mathbb{P}(\Phi(U_1, \ldots, U_T) \in A \mid T < S) \tag{6.2}$$

which is different from, though as $S \to \infty$ converges to, (6.1). This is the so-called *impatient-user bias*. The CFTP algorithm also permits the construction of a joint realization (η, ξ) with marginals (6.1) and (6.2) such that $T < S$ implies $\eta = \xi$. In fact, as

pointed in Proposition 6.2 of Fill (1998)

$$\sup_A |\mathbb{P}(\Phi(U_1, \ldots, U_T) \in A \mid T < S) - \mu(A)| \leqslant \frac{\mathbb{P}[T > S]}{1 - \mathbb{P}[T > S]}. \tag{6.3}$$

In our algorithm T is determined by the number of uniform random variables necessary to construct the clan of the observed region Λ.

When the possible sizes of the individuals γ form an unbounded set, for instance for the Peierls contours of the Ising model, practical limitations prevent the inclusion of all possible sizes in the simulation. In fact the mere enumeration of the possible contours is beyond reach when more than a few dozens of links are involved. This is tantamount to a "space impatient-user bias": the user is forced to do a space cut-off that produces a bias, even when the actual probability for a cut event to take place is tiny. In mathematical terms, one actually samples from the conditioned measure

$$\mathbb{P}(\Phi(U_1, \ldots, U_T) \in A \mid \{K < k\} \cap \{T < S\}), \tag{6.4}$$

where $K =$ "maximum perimeter of bases of cylinders in the clan" ($k = 30$, for instance). In fact, our approach also admits a joint realization (η, ξ) with marginal distributions (6.1) and (6.4) such that $\eta = \xi$ if $K < k$ and $T < S$, and such that $\mathbb{P}(\{K \geqslant k\} \cap \{T > S\})$ goes to zero exponentially fast in S and in the cut-off of the length of the contours (30 in our example). Slightly more precisely, a bound like (6.3) holds with

$$\mathbb{P}(\{K \geqslant k\} \cap \{T > S\}) \leqslant \mathrm{O}\left(\alpha^T \times \sup_x \pi_x(K > k)\right). \tag{6.5}$$

This follows from the subcriticality of the majorizing branching process. For the Ising model, for instance, $\pi_x(K > k) = \mathrm{O}(\mathrm{e}^{-\beta k})$.

7. Conclusion

Our algorithm offers an approach to perfect simulations of processes with infinite state space. The fact that there is no coupling between different initial conditions, makes it a flexible tool for processes with a large state space. No "sandwiching processes" need to be followed; the free process is a natural "dominating process" in our setting. In addition, our algorithm is backed by a rather detailed theory that allows the estimation of various properties of the resulting measure, as well as possible errors. In particular, our approach is not free from the "impatient-user bias", but the resulting error is relatively straightforward to control.

A noteworthy feature of our approach is that the perfect simulation stage applies, in fact, to the free process. Interacting processes are then obtained by a deterministic "cleaning". As a consequence our scheme allows the simultaneous simulation of all processes absolutely continuous with respect to the same free process. This coupled construction could be potentially useful, for instance to establish comparison criteria.

The algorithm admits a further generalization more or less immediate that has not been pursued here: it can be applied to processes with variable death rate that,

however, must be uniformly bounded from below by 1, say. The dependences in the birth and death-rates induce definitions of incompatibility and respective parameters α (cf. (4.10), (4.3)). A construction analogous to the one described in this paper can be performed, but with a thinning algorithm that takes also into account the variable deathrates.

In this work, the advantages of the approach have been exploited only at a theoretical level, where it has led to a new treatment of systems with exclusions and to better estimates of regions of existence of a number of processes. Berthelsen and Møller (2001) compared it to the dominated CFTP introduced by Kendall and Møller (2000). Based on simulation results, the authors show that the dominated CFTP is better than the algorithm based on the clan of ancestors in the particular case of a Strauss process (see Eq. (2.6)) defined on a unit square with $e_1^\beta = 100$ and $e_2^\beta = 0$ (the so-called hard-core process), 0.5 and 1 (a Poisson processes with rate 100). This is obviously the case from the description of the processes since the backward construction of our algorithm stops when the dominated Poisson process regenerates and usually the coupling of CFTP is achieved before it in the finite case. However, it should be noticed that the algorithm based on the clan of ancestors was designed for sampling the infinite-volume process viewed in a finite window. This seems to be a much more interesting and challenging problem which has been studied by Máric (2002) for the specific case of one-dimensional loss networks with bounded calls. No comparison was made to other perfect simulation schemes.

Finally, we hope that a suitable combination of our ideas with some rejection sampling scheme could yield a version free of the user-impatience bias.

Acknowledgements

Appendix A. Proof of Theorems

Proof of Theorem 1. We need to show that η_t^ζ has generator (3.4). Denote $\eta_t = \eta_t^\zeta$ and $\mathbf{K}[0,t]$ the set of kept cylinders born at time zero or after time zero (this includes the cylinders induced by the initial configuration) and for F a function depending on

individuals intersecting a region with finite total rate, write

$$[F(\eta_{t+h}) - F(\eta_t)]$$

$$= \sum_{C \in \mathbf{K}[0,t+h]} 1\{\text{birth}(C) \in [t, t+h]\}[F(\eta_t + \delta_{\text{basis}(C)}) - F(\eta_t)]$$

$$+ \sum_{C \in \mathbf{K}[0,t]} 1\{\text{life}(C) \ni t, \text{life}(C) \not\ni t+h\}[F(\eta_t - \delta_{\text{basis}(C)}) - F(\eta_t)]$$

$$+ \{\text{other things}\}, \tag{A.1}$$

where {other things} refer to events with more than one Poisson mark in the time interval $[t, t+h]$ for the contours in the (finite) support of F. Since the total rate of the Poisson marks in this set is finite, the event {other things} has a probability of order $(hm(\text{Supp}(F), \mathbf{G}))^2$, where m is defined in (4.7). We have

$$\sum_{C \in \mathbf{C}} 1\{\text{birth}(C) \in [t, t+h]\} 1\{C \in \mathbf{K}[0, t+h]\}[F(\eta_t + \delta_{\text{basis}(C)}) - F(\eta_t)]$$

$$= \sum_{C \in \mathbf{C}} 1\{\text{birth}(C) \in [t, t+h]\} 1\{\text{flag}(C) < M(\text{basis}(C)|\eta_t)\}$$

$$\times [F(\eta_t + \delta_{\text{basis}(C)}) - F(\eta_t)]. \tag{A.2}$$

To compute the second term of (A.1), observe that $\text{life}(C)$ is independent of $\text{birth}(C)$ and both the event $\{C \in \mathbf{K}[0, t]\}$ and η_t are $\mathscr{F}_t$-measurable. Here $\mathscr{F}_t$ is the σ-algebra generated by the births and deaths occurred before t. Hence

$$\mathbb{P}(\text{life}(C) \ni t, \text{life}(C) \not\ni t+h | \mathscr{F}_t)$$

$$= \mathbb{P}(\text{life}(C) \not\ni t+h \,|\, \text{life}(C) \ni t) 1\{\text{life}(C) \ni t\}$$

and

$$\mathbb{E}\left[\sum_C 1\{C \in \mathbf{K}[0,t]\} 1\{\text{life}(C) \ni t, \text{life}(C) \not\ni t+h\}[F(\eta_t - \delta_{\text{basis}(C)}) - F(\eta_t)] \right]$$

$$= \mathbb{E}\left[\sum_C \mathbb{P}(\text{life}(C) \not\ni t+h \,|\, \text{life}(C) \ni t) 1\{C \in \mathbf{K}[0,t], \text{life}(C) \ni t\} \right.$$

$$\left. \times [F(\eta_t - \delta_{\text{basis}(C)}) - F(\eta_t)] \right]. \tag{A.3}$$

Since $\text{life}(C)$ is exponentially distributed with mean 1,

$$\mathbb{P}(\text{life}(C) \not\ni t+h \,|\, \text{life}(C) \ni t) = h + \mathrm{o}(h). \tag{A.4}$$

Taking the expectation of (A.1) and substituting (A.2)–(A.4) we get

$$
\mathbb{E}[F(\eta_{t+h}) - F(\eta_t)]
$$

$$
= h \int_{\mathbf{G}} \nu(\mathrm{d}\gamma) \mathbb{E}(M(\gamma \,|\, \eta_t)[F(\eta_t + \delta_\gamma) - F(\eta_t)]) + \mathrm{o}(h)
$$

$$
+ h \sum_{\gamma:\eta_t(ga)>0} \mathbb{E}(\eta_t(\gamma)[F(\eta_t - \delta_\gamma) - F(\eta_t)]) + \mathrm{o}(h) \tag{A.5}
$$

which dividing by h and taking limit gives

$$
\frac{\mathrm{d}\mathbb{E}F(\eta_t^\zeta)}{\mathrm{d}t} = A\mathbb{E}F(\eta_t^\zeta). \qquad \square \tag{A.6}
$$

Proof of Theorem 2. Define

$$
\mathbf{A}_t = \{C' \in \mathbf{A}^{\Lambda,0}: 0 > \mathrm{birth}(C') > -t\} = \mathbf{A}^{\Lambda,0} \cap \mathbf{C}[-t,0], \tag{A.7}
$$

that is, the set of cylinders in $\mathbf{A}^{\Lambda,0}$ with birth time posterior to $-t$. It suffices to prove that the process so defined satisfies the evolution equation (5.3).

The inclusion of a new cylinder in the time interval $[t, t+h]$ depends on the existence of a birth Poisson mark in $[-t-h, -t]$ whose corresponding cylinder is incompatible with some $C' \in \mathbf{A}_t$. That is, if C is a cylinder with $I(\mathrm{basis}(C'), \mathrm{basis}(C)) = 1$ for some $C' \in \mathbf{A}_t$,

$$
\mathbb{P}(\mathbf{A}_{t+h} = \tilde{\mathbf{A}} \cup C \,|\, \mathbf{A}_t = \tilde{\mathbf{A}}, \mathbf{A}_{t'} = \tilde{\mathbf{A}}_{t'}, t' \in [0,t))
$$

$$
= \mathbb{P}\{C \in \mathbf{C}: \mathrm{birth}(C) \in [-t-h, -t], \mathrm{death}(C) > t - \mathrm{TI}(\tilde{\mathbf{A}}, \Lambda, \mathrm{basis}(C))\} + \mathrm{o}(h).
$$

The remainder $\mathrm{o}(h)$ is the correction related to the probability that C is not the only *relevant* cylinder born in $[-t-h, -t]$. Hence

$$
\mathrm{o}(h) \leqslant \left(h \sum_{C \in \mathbf{A}_t} m(\mathrm{basis}(C), \mathbf{G}) \right)^2 \leqslant h^2 |\mathbf{A}_t|^2 \alpha^2,
$$

where m is defined in (4.7) and $|\mathbf{A}_t|$ stands for $\sum_{C \in \mathbf{A}_t} q(\mathrm{basis}(C))$, where q is the measure used to define α in (4.10). Since the birth time is independent of the lifetime which is exponentially distributed with rate one,

$$
\mathbb{P}(\mathbf{A}_{t+h} = \tilde{\mathbf{A}} \cup C \,|\, \mathbf{A}_t = \tilde{\mathbf{A}}, \mathbf{A}_{t'} = \tilde{\mathbf{A}}_{t'}, t' \in [0,t))
$$

$$
= \mathbb{P}\{C \in \mathbf{C}: \mathrm{birth}(C) \in [-t-h, -t]\}
$$

$$
\times \mathbb{P}(\mathrm{life}(C) > t - \mathrm{TI}(\tilde{\mathbf{A}}, \Lambda, \mathrm{basis}(C))) + \mathrm{o}(h)
$$

$$
= h f(\mathrm{basis}(C)) \, \mathrm{e}^{-t + \mathrm{TI}(\tilde{\mathbf{A}}, \Lambda, \mathrm{basis}(C))} + \mathrm{o}(h). \tag{A.8}
$$

This implies that when the configuration at time t^- is $\tilde{\mathbf{A}}$, a new cylinder with basis γ is included in $\mathbf{A}_t(\Upsilon)$ at rate

$$
f(\gamma) \, \mathrm{e}^{-t + \mathrm{TI}(\tilde{\mathbf{A}}, \Lambda, \gamma)}. \tag{A.9}
$$

From ([A.8](#)), as in the computation of the forward Kolmogorov equations, we get ([5.3](#)). This equation characterizes the law of the process $\mathbf{A}_t(\Upsilon)$ as a non-homogeneous Markov process. □

References

Baddeley, A.J., van Lieshout, M.N.M., 1995. Area-interaction point processes. Ann. Inst. Statist. Math. 47 (4), 601–619.

Berthelsen, K.K., Møller, J., 2001. Spatial jump processes and perfect simulation. Preprint.

Brydges, D.C., 1986. A short course on cluster expansions. In: Phénomènes critiques, systèmes aléatoires, théories de gauge, Part I, II (Les Houches, 1984). North-Holland, Amsterdam, 1986, pp. 129–183.

Cai, H., 1999. Exact sampling using auxiliary variables. ASA Proc. (Statist. Comput. Section), to appear.

Comets, F., Fernandez, R., Ferrari, P.A., 2002. Processes with long memory: regenerative construction and perfect simulation. Ann. Appl. Probab., arXiv:math.PR/0009204, to appear.

Corcoran, J.N., Tweedie, R.L., 2001. Perfect sampling of ergodic Harris chains. Ann. Appl. Probab. 11 (2), 438–451.

Dobrushin, R.L., 1996. Perturbation methods of the theory of Gibbsian fields. In: Lectures on Probability Theory and Statistics (Saint-Flour,1994). Springer, Berlin, pp. 1–66.

Fernández, R., Fröhlich, J., Sokal, A.D., 1992. Random Walks, Critical Phenomena, and Triviality in Quantum Field Theory. Springer, Berlin, Heidelberg, New York.

Fernández, R., Ferrari, P.A., Garcia, N.L., 1998. Measures on contour, polymer or animal models. A probabilistic approach. Markov Process. Related Fields 4 (4), 479–497. I Brazilian School in Probability (Rio de Janeiro, 1997).

Fernández, R., Ferrari, P.A., Garcia, N.L., 2001. Loss network representation of Peierls contours. Ann. Probab. 29 (2), 902–937.

Ferrari, P.A., 1990. Ergodicity for spin systems with stirrings. Ann. Probab. 18 (4), 1523–1538.

Ferrari, P.A., Garcia, N.L., 1998. One-dimensional loss networks and conditioned $M/G/\infty$ queues. J. Appl. Probab. 35 (4), 963–975.

Fill, J.A., 1998. An interruptible algorithm for perfect sampling via Markov chains. Ann. Appl. Probab. 8 (1), 131–162.

Fill, J.A., Huber, M., 2000. The randomness recycler: a new technique for perfect sampling. Preprint. arXiv:math.PR/0009242.

Fill, J.A., Machida, M., Murdoch, D.J., Rosenthal, J.S., 2000. Extension of Fill's perfect rejection sampling algorithm to general chains. Proceedings of the Ninth International Conference "Random Structures and Algorithms" (Poznan, 1999). Random Struct. Algorithms 3–4, 290–316.

Fortuin, C.M., Kasteleyn, P.W., 1972. On the random cluster model. I. Introduction and relation with other models. Physica 57, 536–564.

Foss, S.G., Tweedie, R.L., 1998. Perfect simulation and backward coupling. Stochastic Models 14, 187–203.

Garcia, N.L., 2000. Perfect simulation of spatial processes. Resenhas IME-USP 4 (3), 281–324.

Gawędzki, K., Kotecký, R., Kupiainen, A., 1987. Coarse graining approach to first-order phase transitions. J. Statist. Phys. 47, 701–724.

Grimmett, G., 1995. The stochastic random-cluster process and the uniqueness of random-cluster measures. Ann. Probab. 23 (4), 1461–1510.

Häggström, O., Steif, J., 2000. Propp–Wilson algorithms and finitary codings for high noise random fields. Combin. Probab. Comput. 9 (5), 425–439.

Häggström, O., van Lieshout, M.N.M., Møller, J., 1999. Characterization results and Markov chain Monte Carlo algorithms including exact simulation for some spatial point processes. Bernoulli 5 (4), 641–658.

Hobert, J.P., Robert, C.P., 2000. Moralizing perfect sampling. Preprint.

Kelly, F.P., 1991. Loss networks. Ann. Appl. Probab. 1 (3), 319–378.

Kelly, F.P., Ripley, B.D., 1976. A note on the Strauss' model for clustering. Biometrika 63, 357–360.

Kendall, W.S., 1997. On some weighted Boolean models. In: Jeulin, D. (Ed.), Proceedings of the International Symposium on Advances in Theory and Applications of Random Sets (Fontainebleau, 1996). World Scientific Publishing, River Edge, NJ, pp. 105–120.

Kendall, W.S., 1998. Perfect simulation for the area-interaction point process. In: Accardi, L., Heyde, C.C. (Eds.), Probability Towards 2000. Springer, Berlin, pp. 218–234.

Kendall, W.S., Møller, J., 2000. Perfect simulation using dominating processes on ordered spaces, with applications to locally stable point processes. Adv. Appl. Probab. 32 (3), 844–865.

Kotecký, R., Preiss, D., 1986. Cluster expansion for abstract polymer models. Comm. Math. Phys. 103 (3), 491–498.

Máric, N., 2002. Perfect simulation for a continuous one-dimensional loss network. Master's Thesis, IMECC/UNICAMP. Available at http://www.ime.unicamp.br/rel_pesq/2002/rp18-02.html.

Møller, J., 2000. A review on perfect simulation in stochastic geometry. IMS Lecture Notes, Monograph Series, Hayward, CA.

Olivieri, E., Picco, P., 1990. Cluster expansion for D-dimensional lattice systems and finite volume factorization properties. J. Statist. Phys. 59, 221–256.

Propp, J.G., Wilson, D.B., 1996. Exact sampling with coupled Markov chains and applications to statistical mechanics. Proceedings of the Seventh International Conference on Random Structures and Algorithms (Atlanta, GA, 1995), Vol. 9. pp. 223–252.

Strauss, D.J., 1975. A model for clustering. Biometrika 62 (2), 467–475.

Thönnes, E., 1999. Perfect simulation of some point processes for the impatient user. Adv. Appl. Probab. 31, 69–87.

van den Berg, J., 1993. A uniqueness condition for Gibbs measures, with application to the 2-dimensional Ising antiferromagnet. Comm. Math. Phys. 152 (1), 161–166.

van den Berg, J., Maes, C., 1994. Disagreement percolation in the study of Markov fields. Ann. Probab. 22 (2), 749–763.

van den Berg, J., Steif, J.E., 1999. On the existence and nonexistence of finitary codings for a class of random fields. Ann. Probab. 27 (3), 1501–1522.