




ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Ann. I. H. Poincaré – PR 40 (2004) 141–152



www.elsevier.com/locate/anihpb

Poisson trees, succession lines and coalescing random walks

P.A. Ferrari^a, C. Landim^{b,c}, H. Thorisson^d

^a *IME USP, Caixa Postal 66281, 05311-970 São Paulo, Brazil*

^b *IMPA, Estrada Dona Castorina 110, CEP 22460 Rio de Janeiro, Brazil*

^c *CNRS UMR 6085, Université de Rouen, 76128 Mont Saint Aignan, France*

^d *ICE-UICE-SI, Science Institute, University of Iceland (Háskóli Islands), 107 Reykjavík, Iceland*

Received 16 October 2002; received in revised form 18 July 2003

Abstract

We give a deterministic algorithm to construct a graph with no loops (a tree or a forest) whose vertices are the points of a d -dimensional stationary Poisson process $S \subset \mathbb{R}^d$. The algorithm is independent of the origin of coordinates. We show that (1) the graph has one topological end – that is, from any point there is exactly one infinite self-avoiding path; (2) the graph has a unique connected component if $d = 2$ and $d = 3$ (a tree) and it has infinitely many components if $d \geq 4$ (a forest); (3) in $d = 2$ and $d = 3$ we construct a bijection between the points of the Poisson process and \mathbb{Z} using the preorder-traversal algorithm. To construct the graph we interpret each point in S as a space-time point $(x, r) \in \mathbb{R}^{d-1} \times \mathbb{R}$. Then a $(d - 1)$ -dimensional random walk in continuous time continuous space starts at site x at time r . The first jump of the walk is to point x' , at time $r' > r$, $(x', r') \in S$, where r' is the minimal time after r such that $|x - x'| < 1$. All the walks jumping to x' at time r' coalesce with the one starting at (x', r') . Calling $(x', r') = \alpha(x, r)$, the graph has vertex set S and edges $\{(s, \alpha(s)), s \in S\}$. This enables us to shift the origin of $S^\circ = S \cup \{0\}$ (the Palm version of S) to another point in such a way that the distribution of S° does not change (to any point if $d = 2$ and $d = 3$; point-stationarity).

© 2004 Elsevier SAS. All rights reserved.

Résumé

Nous présentons un algorithme déterministe pour construire un graphe sans boucles (un arbre ou une forêt) dont les sommets sont les points d'un processus de Poisson stationnaire $S \subset \mathbb{R}^d$. L'algorithme est indépendant de l'origine des coordonnées. Nous démontrons que (1) le graphe a une fin topologique – c'est à dire, que de n'importe quel point il existe exactement un seul chemin sans intersections; (2) le graphe a une seule composante connexe pour $d = 2$ et $d = 3$ (un arbre) et il a un nombre infini de composantes pour $d \geq 4$ (une forêt); (3) pour $d = 2$ et $d = 3$, nous construisons une bijection entre les points du processus de Poisson et \mathbb{Z} en utilisant l'algorithme "preorder-traversal".

Pour construire le graphe, nous interprétons chaque point de S comme un point spatio-temporel $(x, r) \in \mathbb{R}^{d-1} \times \mathbb{R}$. Une marche aléatoire $(d - 1)$ -dimensionnelle en temps continu commence du site x à l'instant r . Le premier saut de la marche est vers le point x' à l'instant $r' > r$, $(x', r') \in S$ où r' est le premier instant après r tel que $|x - x'| < 1$. Toutes les marches qui sautent vers x' à l'instant r' s'unissent à la marche débutant en (x', r') pour devenir une seule marche aléatoire. Si $(x', r') = \alpha(x, r)$, l'ensemble S représente les sommets du graphe et l'ensemble $\{(s, \alpha(s)), s \in S\}$ les arêtes. Ceci nous permet

E-mail addresses: pablo@ime.usp.br (P.A. Ferrari), landim@impa.br (C. Landim), hermann@hi.is (H. Thorisson).

URLs: <http://www.ime.usp.br/~pablo> (P.A. Ferrari), <http://www.impa.br/Pesquisadores/Claudio/> (C. Landim), <http://www.hi.is/~hermann/> (H. Thorisson).

0246-0203/\$ – see front matter © 2004 Elsevier SAS. All rights reserved.
doi:10.1016/j.anihpb.2003.12.001

de translater l'origine de $S^\circ = S \cup \{0\}$ (la version de Palm de S) vers un autre point de manière à ce que la distribution de S° reste inchangée.

© 2004 Elsevier SAS. All rights reserved.

MSC: 60K35

Keywords: Poisson processes; Random trees; Palm measure; Coalescing random walks; Point-stationarity; Palm theory

1. Introduction

Let S° be the Palm version of a stationary Poisson process S in \mathbb{R}^d , that is, S° is a random set with the same distribution as $S \cup \{0\}$. In the open problem session at the Brazilian School of Probability in 2001 the third author posed the following three problems. Let S be the support of S and S° the support of S° .

Problem 1.1. When $d > 1$, is there some non-randomized way of shifting the origin of S° from the point at the origin to another point $X \in S^\circ$ so that the distribution of S° as seen from this point remains the same? More precisely, is there a map $\pi : S^\circ \rightarrow \mathbb{R}^d \setminus \{0\}$ such that with $X = \pi(S^\circ) \in S^\circ \setminus \{0\}$ and, with $S^\circ - X = \{s - X : s \in S^\circ\}$,

$$S^\circ - X = S^\circ \quad \text{in distribution?} \quad (1.1)$$

This is clearly possible when $d = 1$, since then the intervals between points are i.i.d. exponential and remain so when the origin is shifted to the n th point on the right (or on the left) of the point at the origin. When $d > 1$, an X satisfying (1.1) – and with $P(X \neq 0)$ arbitrarily close to 1 – exists if external randomization is allowed. This is shown by Thorisson [13,14, Chapter 9] who also proved that (1.1) holds if the *point shift* map $\theta_\pi : S^\circ \rightarrow S^\circ$ defined by

$$\theta_\pi(S^\circ) := S^\circ - \pi(S^\circ) = S^\circ - X$$

is *bijective*.

Problem 1.2. Does there exist a family of maps $\pi_n : S^\circ \rightarrow \mathbb{R} \setminus \{0\}$ such that, defining $X_n := \pi_n(S^\circ)$, $n \in \mathbb{Z}$, we have $X_n \neq X_{n'}$ if $n \neq n'$ and

$$S^\circ = \{X_n : n \in \mathbb{Z}\} \quad (1.2)$$

(that is, a labelling of all the points of S°) and the maps $\theta_{\pi_n} : S^\circ \rightarrow S^\circ$ defined by

$$\theta_{\pi_n}(S^\circ) := S^\circ - X_n$$

are bijective? As we just have seen, this would imply that $S^\circ - X_n = S^\circ$, in distribution.

Problem 1.3. In the references mentioned before, Thorisson defined “point-stationarity” of a point process as distributional invariance under bijective point-shifts “against any independent stationary background”; this concept is shown to be the characterizing property of the Palm version of any stationary point process S in \mathbb{R}^d . A natural question is whether “against any independent stationary background” can be removed from the definition. In other words, can the definition of “point-stationarity” be reduced to distributional invariance under non-randomized bijective point-shifts?

Olle Häggström has given an example of a non-randomized bijection as in Problem 1.1: let X be the closest point to the point at the origin if the point at the origin is also the closest point to that point; otherwise let X be 0. But is there a strictly non-zero X ? Dana Randall and the first author proposed the following map in the same vein: in the first step “marry” each point s to its closest point s' if s is the closest point to s' . Call $M(S)$ the set of points in S married using this procedure and $M_1 = M(S)$. The set of points married in the n th step is

$M_n = M(S \setminus (M_1 \cup \dots \cup M_{n-1}))$. Every point will eventually get married, that is $\bigcup_n M_n = S$, because the Poisson process has no descending chains, as proved by Häggström and Meester [8]; this has been observed by Holroyd and Peres [9]. In this case, the map $s \mapsto (\text{spouse of } s)$ for $s = 0$ is such a non-zero X for all d .

We provide another solution to 1.1 for all d and show that the answer to 1.2 is “yes” when $d = 2$ and $d = 3$. We give a partial response to 1.3.

We show that in $d = 2$ and $d = 3$ it is possible to join the points of S° in an origin-independent way into a connected tree with finite branches. For this tree every pair of vertices have an ancestor in common and every vertex has a finite number of descendants. The resulting tree is called *Poisson tree*. We order sets of sisters using the first spacial coordinate and associate to each vertex a semi-infinite sequence indicating the sister-order of her ancestors. Then we order these sequences lexicographically from the past to get a “unique infinite succession line” of vertices. The total order so obtained corresponds to the *preorder traversal* algorithm used in computer science.

Let $\pi_0(S^\circ) := 0$ and inductively $\pi_n(S^\circ) := X_n$, where X_n is the successor of X_{n-1} for $n \geq 1$ and the predecessor of X_{n+1} for $n \leq 1$ in the total order just described. Let $\theta_{\pi_n} : S^0 \mapsto S^0 - X_n$. It is clear that each point is the successor of its predecessor and more generally the n th successor of its n th predecessor and vice versa:

$$\pi_{-n}(S^\circ - \pi_n(S^\circ)) = -\pi_{-n}(S^\circ)$$

for $n \geq 0$, which implies $\theta_{\pi_n}^{-1} = \theta_{\pi_{-n}}$ for all n . Thus, for all $n \in \mathbb{Z}$, θ_{π_n} is a bijection as required in 1.2. Notice also that $\theta_{\pi_n} = \theta_{\pi_1}^n$ the n th iteration of θ_{π_1} . In $d \geq 4$ we construct a “forest”, infinitely many connected trees with finite branches. Each tree can be ordered using the preorder traversal algorithm and the map $s \mapsto (\text{successor of } s)$ is a solution to Problem 1.1 in any $d \geq 1$.

In $d = 2$ (one dimension for space and the other for time) our construction is a continuous space-time analogous of a discrete space-time system of coalescing random walks. Both the system of coalescing random walks and the Poisson tree converge to the so called “Brownian web”, a system of one-dimensional coalescing Brownian motions starting at every space-time point in \mathbb{R}^2 . The random walk convergence and different properties of the web have been studied by Arratia [2,3] and Toth and Werner [15]. Ferrari, Fontes and Wu [5] show that the Poisson tree converges to the Brownian web in the sense proposed by Fontes, Isopi, Newman and Ravishankar [6].

Gangopadhyay, Roy and Sarkar [7] proposed a system of coalescing random walks based on a Bernoulli product measure in \mathbb{Z}^2 and show that it produces a connected tree. Their motivation is to provide a model for drainage networks; see the book of Rodriguez-Iturbe and Rinaldo [12].

A tree having S as vertex set is the *minimal spanning tree* constructed as follows. Choose arbitrarily an initial point $s \in S$. Let (V_n, E_n) be the set of vertices and edges chosen up to the n th iteration; $(V_0, E_0) = (\{s\}, \emptyset)$. The $(n + 1)$ th point s_{n+1} is the point in the complement of V_n that is closest to V_n . The $(n + 1)$ th edge is the pair (s', s_{n+1}) , $s' \in V_n$ realizing the distance between V_n and s_{n+1} . Alexander [1] proved that in $d = 2$ the construction is independent of the initial point, that $V_n \rightarrow S$ and that the resulting tree has all branches finite. In this terminology, the tree has one topological end – from each point there is only one infinite self-avoiding path. It is believed that for $d \leq 8$ this tree has a unique connected component; Newman and Stein [10,11].

In Section 2 we relate trees and succession lines. In Section 3 we construct the Poisson tree and state that in $d = 2$ and $d = 3$ the tree is connected and has finite branches (Theorem 3.1). In Section 4 we introduce a graphical construction of a system of coalescing random walks, associate it to the tree and prove Theorem 4.3 – from where Theorem 3.1 follows. In that section we also show that the system of coalescing random walks is ergodic and converges exponentially fast to the unique invariant measure. In Section 5 we give a limited reply to Problem 1.3. In Section 6 we make some final remarks and state related open problems.

2. Trees and succession lines

Let \mathcal{T} be an oriented graph with no loops (a tree or a forest) such that each vertex s of \mathcal{T} is the startpoint of exactly one outgoing edge and the endpoint of either none or a finite number of ingoing edges. The endpoint of

the outgoing edge is called the *mother* of s , while the startpoints of the ingoing edges are the *daughters* of s . If two vertices are daughters of the same vertex, they are *sisters* of each other. We order groups of sisters according with some spacial property like first coordinate, distance to the mother, etc.; we denote $s < s'$ and say that s is older than s' in the adopted order. This order is not necessarily coherent with the partial order induced by the ancestor-descendent relation.

Call $\alpha(s, T)$ the mother of s . Let $\alpha^0(s, T) = s$, and iteratively, for $n \geq 1$, $\alpha^n(s, T) = \alpha(\alpha^{n-1}(s, T), T)$ the n th ancestor of s . Let

$$\begin{aligned} D^1(s, T) &:= \{s' \in S : \alpha(s', T) = s\}, \\ D^n(s, T) &:= \{s' \in S : \alpha(s', T) \in D^{n-1}(s, T)\}, \\ D(s, T) &:= \bigcup_{n \geq 0} D^n(s, T) \end{aligned} \tag{2.1}$$

be respectively the first generation, the n th generation and the set of all descendents of s ; we call $D(s, T)$ the *branch* of s and say that T has finite branches if $D(s, T)$ is finite for all s . We say that two vertices s and s' are *connected* if they have an ancestor in common: there exist nonnegative integers n and m such that $\alpha^n(s, T) = \alpha^m(s', T)$. This defines an equivalence relation in S ; the equivalence classes are called *connected components*.

Let T be a connected tree with finite branches. Let $\sigma(s) = 1$ if s is the eldest among her sisters, $\sigma(s) = 2$ if s is the second sister, and so on. We associate to each vertex s the sequence of relative sister-order of its ancestors: let $\sigma_i(s) := \sigma(\alpha^i(s))$, $i \geq 0$. If $s'' = \alpha^i(s) = \alpha^j(s')$ is the closest common ancestor of s and s' , for nonnegative i, j , then they can be lexicographically ordered using the non-common part of the sequences: we say that s *precedes* s' if $(\sigma_i(s), \sigma_{i-1}(s), \dots, \sigma_1(s), \sigma_0(s))$ is lexicographically before $(\sigma_j(s'), \sigma_{j-1}(s'), \dots, \sigma_1(s'), \sigma_0(s'))$. That is, if $\sigma_{i-1}(s) < \sigma_{j-1}(s')$, with the convention $\sigma_{-1}(s) = 0$.

For $s \in T$, define the *successor* of s as $s' \in T$ if s precedes s' and there is no vertex preceded by s and preceding s' . Conversely s is the *predecessor* of s' if and only if s' is the successor of s .

The successor of a vertex can be found using the following algorithm. If the vertex has a daughter, choose the eldest daughter. If it does not have a daughter but has a younger sister, choose the eldest among its younger sisters. If it does not have a daughter and not a younger sister, move up the tree until you hit the first point that has a younger sister and choose the eldest among its younger sisters. This requires that every vertex has an ancestor with a younger sister.

The predecessor vertex can be found with this algorithm. If the vertex has an elder sister, choose the youngest among her elder sisters and then move from her down the tree choosing the youngest daughter in each step until you come to a point with no daughter; this will be the predecessor. This requires that the branch of the mother of each individual is finite.

We say that there is a *succession line* from s to s' if there exists a finite sequence of vertices $s = s_0, \dots, s_k = s'$ such that $s_{\ell-1}$ is successor of s_ℓ for $\ell = 1, \dots, k$. We say that the tree has an *infinite succession line* if every vertex has a predecessor and a successor and that it has a *unique infinite succession line* if furthermore for every couple of vertices s, s' there is a succession line either from s to s' or from s' to s . The following lemma follows from the definitions.

Lemma 2.1. *If $T = (S, E)$ has a unique infinite succession line, then the map $\pi(s, T) = s'$, the successor of S and $\pi^{-1}(s', T) = s$, the predecessor of s' are well defined (i.e., both the successor and predecessor algorithms find a vertex) and one is the inverse of the other. Furthermore, for all $s \in S$, $S = \{\pi^n(s, T) : n \in \mathbb{Z}\}$.*

The following lemma gives an equivalent condition.

Lemma 2.2. *A tree T has a unique connected component, finite branches and every vertex has a mother and an ancestor with a younger sister if and only if T has a unique infinite succession line.*

Proof. Since every vertex s' has a mother and all branches are finite, the predecessor of s' is one of the vertices of the branch of her mother and can be found in a finite number of steps. The condition that every vertex s has an ancestor with a younger sister guarantees the existence of the successor of s that can also be found in a finite number of steps. Since the tree is connected, two arbitrary vertices s and s' have an ancestor in common, say s'' . The branch of s'' is finite, by hypothesis; say it has n vertices. Call $s''_0 = s''$, s''_k the successor of s''_{k-1} , $k = 1, \dots, n$. Then the branch of s'' is the same as $\{s''_0, \dots, s''_n\}$. In particular it contains s and s' and a succession line either from s to s' or from s' to s . The converse statement is immediate. \square

3. Poisson tree

Let $d \geq 2$ and S be a locally finite configuration of points in \mathbb{R}^d . For each $s = (s_1, \dots, s_d) \in \mathbb{R}^d$ call its first $d - 1$ coordinates $x(s) = (s_1, \dots, s_{d-1}) \in \mathbb{R}^{d-1}$ and the remaining coordinate $r(s) = s_d$. In this way $s = (x(s), r(s))$; $x(s)$ is interpreted as the *space coordinate* and $r(s)$ the *time coordinate* of s .

For each $x \in \mathbb{R}^{d-1}$ let $B(x) = \{x' \in \mathbb{R}^{d-1}; |x' - x| \leq 1\}$ be the $(d - 1)$ -dimensional Euclidean ball of radius 1 centered at x . For each $(x, r) \in S$ call $\{(x', r): x' \in B(x)\}$ (the $(d - 1)$ -dimensional disk centered at (x, r) perpendicular to the d th axis) the *obstacle* associated to (x, r) . The set of obstacles is given by

$$\mathcal{O}(S) = \bigcup_{(x,r) \in S} \{(x', r): x' \in B(x)\}.$$

Think that each point $(x, r) \in S$ emits a laser ray in the positive d th coordinate that is stopped by the obstacles. The first obstacle hit by the ray of (x, r) has second coordinate

$$\tau((x, r), S) := \inf\{t > r: (x, t) \cap \mathcal{O}(S) \neq \emptyset\} \tag{3.1}$$

with center

$$\alpha((x, r), S) := (x', r') \in S \quad \text{if } \tau(s, S) = r' \tag{3.2}$$

which is called the *mother* of (x, r) . Reciprocally, s is a *daughter* of $\alpha(s, S)$. The above objects are well defined for S if $\tau(s, S) < \infty$ for all $s \in S$ and if no point has two mothers. In this case let $\mathcal{T}(S) = (S, E(S))$ be the random directed graph with vertices S and edges $E(S) = \{(s, \alpha(s, S)): s \in S\}$. This graph has no loops, hence it is an oriented tree. Notice that $\alpha(s, S)$ coincides with $\alpha(s, \mathcal{T}(S))$, in the notation of the previous section.

Theorem 3.1. *Let S be the realization of a homogeneous d -dimensional Poisson process and S° its Palm version. Then, for $\mathcal{T} = \mathcal{T}(S)$ and $\mathcal{T} = \mathcal{T}(S^\circ)$ it holds S -a.s. and S° -a.s.:*

- (a) \mathcal{T} is well defined.
- (b) In $d = 2, 3$, \mathcal{T} has a unique connected component.
- (c) In $d \geq 4$, \mathcal{T} has infinitely many connected components.
- (d) All branches of \mathcal{T} are finite.
- (e) Every vertex has a mother.
- (f) Using the order of the first coordinate in \mathbb{R}^d , each vertex has an ancestor with a younger sister.

Items (b) and (d) to (f) are necessary to construct the maps π_n of Problem 1.2. To prove (a) it suffices to see that for almost all realizations of S and S° , (3.2) has a unique solution; that is, each point has a unique mother. This is clear for a Poisson process and its Palm version. The proofs of (b) to (e) are based on a particle system in \mathbb{R}^{d-1} studied in next section; they are direct consequences of Theorem 4.3.

4. Coalescing random walks

Let S be a point configuration satisfying that every point has exactly one mother. Write $\alpha(s)$ instead of $\alpha(s, S)$ and recall the notation introduced in Section 2. For $s \in \mathbb{R}^d$, let

$$\tau^n(s) := r(\alpha^n(s)); \tag{4.1}$$

that is, $\tau^n(s)$ is the time coordinate of the n th ancestor of s . For each $(x, r) \in \mathbb{R}^{d-1} \times \mathbb{R}$ let $X_r^{(x,r)} = x$ and for $t \geq r$,

$$X_t^{(x,r)} = x(\alpha^n(x, r)), \quad \text{for } t \in [\tau^n(x, r), \tau^{n+1}(x, r)), \quad n \geq 0. \tag{4.2}$$

That is, $X_t^{(x,r)}$ starts at x at time r and remains still between $\tau^{n-1}(x, r)$ and $\tau^n(x, r)$, when it jumps to $x(\alpha^n(s))$, the center of the n th obstacle it meets. In other words, $\tau^n(x, r)$ is the instant of the n th jump of the point that at time r was at position x . The family

$$\mathcal{X}(S) := \{ (X_t^{(x,r)} : t \geq r) : (x, r) \in S \} \tag{4.3}$$

is a deterministic function of S .

Let S be a point configuration of a homogeneous Poisson process of rate λ . Then (4.3) is a family of random processes in the probability space where S is defined; its law corresponds to a system of coalescing random walks with births, whose marginal distribution is described by “each random walk waits an exponentially distributed random time of mean $(V_{d-1}\lambda)^{-1}$ after which it chooses a point uniformly in $B(x)$, the $(d - 1)$ -dimensional Euclidean ball of radius one centered at x , and jumps to it”. Here V_d stands for the volume of the d -dimensional ball of radius one. Particles are created at a Poisson rate λ and the interaction appears when two walks are located at points x and y , $|x - y| < 2$ and a Poisson event appears in $B(x) \cap B(y)$; in this case both walks jump to the same point and coalesce. Disregarding the label of the particles, for $t \in \mathbb{R}$, let

$$\eta_t = \{ X_t^{(x,r)} : (x, r) \in S, r \leq t \}.$$

Here η_t is a discrete subset of \mathbb{R}^{d-1} . The process starting from a fixed configuration $\eta \subset \mathbb{R}^{d-1}$ at fixed time t' is defined by:

$$\eta_t^{\eta, t'} = \{ X_t^{(x,r)} : (x, r) \in S, t' \leq r \leq t \text{ or } x \in \eta, r = t' \}$$

for $t \geq t'$. That is, ignore the starting points with second coordinate less than t' and include the points with first coordinate in η and second coordinate equal to t' .

Proposition 4.1. *Let S be a stationary Poisson process. The process $(\eta_t^{\eta, t'} : t \geq t')$ is Markov with generator defined by*

$$Lf(\eta) = \lambda \int_{\mathbb{R}^{d-1}} [f(\eta \setminus B(u) \cup \{u\}) - f(\eta)] du,$$

for functions $f : \mathcal{S}_{d-1} \rightarrow \mathbb{R}$ depending on bounded regions, where \mathcal{S}_{d-1} is the set of locally finite labeled configurations of \mathbb{R}^{d-1} . Furthermore, η_t is a stationary version of the unique stationary process with generator L . In particular, the marginal distribution of η_t for any given t is the unique invariant measure for the generator L .

Proof. The independence properties of the Poisson process S imply that for any initial configuration η , the process $(\eta_t^{\eta, t'}, t \geq t')$ is Markov and has generator L .

To show that the process is well defined starting at time $-\infty$ we show first that the configuration in finite sets depends only on a finite number of points of S . Let A be a subset of \mathbb{R}^{d-1} with finite Lebesgue measure. The idea

is to prove that for almost all configuration S , $\eta_t \cap \Lambda$ depends only on a finite (but random) subset $S_{\Lambda,t}$ of S . To get the set $S_{\Lambda,t}$ one first translates the set Λ backwards in time hitting points (x, r) . Each time a point (x, r) is hit, Λ is updated to $\Lambda \setminus B(x)$ and (x, r) is included in $S_{\Lambda,t}$. The procedure stops when Λ is updated to the empty set. Then for each point (x, r) already in $S_{\Lambda,t}$ include also all its ancestors with second coordinate less than t : $(x^i, r^i) = \alpha^i(x, r)$ with $r^i \leq t$. We leave to the reader to show that $\eta_t \cap \Lambda$ depends only on $S_{\Lambda,t}$.

For each finite measure set Λ , $S_{\Lambda,t}$ has a finite number of points with probability one. This shows that for any fixed t , as $t' \rightarrow -\infty$ the variable $\eta_{t'}^{\eta_t} \cap \Lambda$ converges almost surely to a random variable $\eta_t \cap \Lambda$. Since the law of $\eta_t \cap \Lambda$ does not depend on t , it is invariant for the process. \square

It follows from this construction that the process converges exponentially fast to equilibrium:

Corollary 4.2. Denote by $\{P_t: t \geq 0\}$ the semigroup associated to the generator L . For every bounded function $f: \mathcal{S}_{d-1} \rightarrow \mathbb{R}$ with finite support, there exists a finite constant $C(f)$ such that

$$\|P_t f - E[f]\|_\infty \leq C(f)e^{-C_d \lambda t}$$

for every $t \geq 0$ and some finite constant C_d depending only on the dimension. In this formula $E[\cdot]$ stands for the expectation with respect to the stationary state.

Proof. Fix two configurations η, ξ and a finite cube Λ . For $t \geq 0$, denote by a_t^Λ all points x in Λ such that (x, r) belongs to S for some $0 \leq r \leq t$ and denote by A_t^Λ the union of all $(d - 1)$ -dimensional balls of radius one with center in a_t^Λ :

$$a_t^\Lambda = \bigcup \{x: (x, r) \in S \text{ for some } 0 \leq r \leq t\},$$

$$A_t^\Lambda = \bigcup_{x \in a_t^\Lambda} B(x).$$

For a configuration ζ , denote by η_t^ζ the state at time t of the process which starts from ζ . By construction $\eta_t^\eta \cap \Lambda = \eta_t^\xi \cap \Lambda$ if A_t^Λ contains Λ .

Divide the cube Λ in small cubes, in such a way that if each small cube contains a point in a_t^Λ then A_t^Λ contains Λ . Denote these cubes by $\{E_i: 1 \leq i \leq M\}$ and notice that the number M of them is equal to $C|\Lambda|$ for some constant C depending only on the dimension. In view of the two previous observations,

$$\mathbb{P}[\eta_t^\eta \cap \Lambda \neq \eta_t^\xi \cap \Lambda] \leq \mathbb{P}[\Lambda \not\subset A_t^\Lambda] \leq \mathbb{P}\left[\bigcup_{i=1}^M \Omega_i(t)\right],$$

where $\Omega_i(t)$ is the event that the cube E_i does not contain points in a_t^Λ . Since a point appears in a set A at rate $\lambda|A|$, $\mathbb{P}[\Omega_i(t)] \leq \exp\{-\lambda|E_i|t\}$ so that

$$\mathbb{P}[\eta_t^\eta \cap \Lambda \neq \eta_t^\xi \cap \Lambda] \leq C_1 |\Lambda| e^{-C_2 \lambda t}$$

for two finite constants depending only on the dimension because $M = C(d)|\Lambda|$ and $|E_i| = C(d)$.

To conclude the proof, it remains to consider a bounded function f with finite support. Denote by Λ a cube which contains its support. $\|P_t(f) - E[f]\|_\infty$ is bounded by

$$\sup_{\eta, \xi} \mathbb{E}[|f(\eta_t^\eta) - \mathbb{E}[f(\eta_t^\xi)]|] \leq 2\|f\|_\infty \mathbb{P}[\eta_t^\eta \cap \Lambda \neq \eta_t^\xi \cap \Lambda],$$

which concludes the proof of the lemma. \square

Theorem 4.3. *Let S be the realization of a homogeneous d -dimensional Poisson process and S° its Palm version. Then, for $\mathcal{X} = \mathcal{X}(S)$ and $\mathcal{X} = \mathcal{X}(S^\circ)$ it holds S -a.s. and S° -a.s.:*

- (a) \mathcal{X} is well defined.
- (b) In $d = 2, 3$, every couple of walkers $X_t^{x,r}, X_t^{x',r'}$ in \mathcal{X} will eventually meet.
- (c) In $d \geq 4$, there are infinitely many walkers in \mathcal{X} that do not meet.
- (d) Each walker alive at time t was born at a finite time before t .
- (e) Every walk will eventually jump.
- (f) Every walk will eventually coalesce with a younger walk (in the order of the first coordinate at coalescence time).

Proof. (a) and (e) follow from the properties of the Poisson process: Each walk will eventually hit a unique obstacle and jump to it.

By symmetry each walk $X_t^{(x,r)}$ is a martingale: Letting \mathcal{F}_t be the σ -algebra generated by $\{s \in S: s_d \leq t\}$,

$$\mathbb{E}(X_t^{(x,r)} | \mathcal{F}_t) = X_t^{(x,r)}, \quad t' \geq t \geq r.$$

(b) In $d = 2$, for $x' < x$, the difference walk $D_t := X_t^{(x,r)} - X_t^{(x',r')} \in \mathbb{R}$, $t \geq \max\{r, r'\}$ is a positive martingale. To check it is enough to show that $\partial_t \mathbb{E}(D_t | X_t^{(x,r)} = a, X_t^{(x',r')} = b) = 0$, which follows from an elementary computation divided in two cases: $a < b - 1 < a + 1 < b$ and $a + 1 < b - 1$. Hence D_t it is recurrent in the sense that $D_t < 1/2$ will occur for an infinite number of times t . Since each time $D_t < 1$ there is a positive probability that $D_{t+1} = 0$, and 0 is an absorbing point for D_t , this implies that (x, r) and (x', r') have an ancestor in common. Since this follows for all points, there is only one tree in $d = 2$.

For $d = 3$, let Y_t and Y'_t be two $(d - 1)$ -dimensional independent random walks with the same marginal distribution as $X_t = X_t^{(x,r)}$ and $X'_t = X_t^{(x',r')}$ respectively. Without loss of generality we can assume $t \geq 0$. The generator \mathcal{L} of the process (Y_t, Y'_t) is given by

$$(\mathcal{L}f)(y, y') = \lambda \int_{B(y)} [f(z, y') - f(y, y')] dz + \lambda \int_{B(y')} [f(y, z) - f(y, y')] dz.$$

Since in dimension 2, $\log |z|$ is a harmonic function, $\mathcal{L} \log |y - y'| = 0$ so that $\log |Y_t - Y'_t|$ is a local martingale. This observation and standard arguments involving the hitting time of two balls centered at the origin and of radius $\underline{R} < |y - y'| < \bar{R}$ permits to show that $Y_t - Y'_t$ is recurrent.

Assume $|X_0 - X'_0| > 2$ and couple $((X_t, X'_t), (Y_t, Y'_t))$ up to the first time T such that $|Y_t - Y'_t| < 2$; for $t \in [0, T]$ we have $X_t = Y_t$ and $X'_t = Y'_t$. Then wait up to the first time $T' > T$ such that either $X_{T'} = X'_{T'}$ or $|X_{T'} - X'_{T'}| > 2$. Since there is a positive probability that $X_{T'} = X'_{T'}$ and the recurrence of $Y_t - Y'_t$ guarantees the existence of infinitely many attempts, $X_t = X'_t$ eventually with probability one.

(c) For $d \geq 4$, consider the process (Y_t, Y'_t) defined in (b). By similar reasons to the ones presented in (b), $|Y_t - Y'_t|^{2-d}$ is a local martingale. This observation and standard arguments involving the hitting time of two balls centered at the origin and of radius $\underline{R} < |y - y'| < \bar{R}$ permits to show that the process $Y_t - Y'_t$ is transient.

Fix $\underline{R} > 2$ and proceed as in the proof of (a) and couple the independent and the interacting walks in such a way that they coincide if the distance is bigger than two. Hence, for initial points y, y' with $|y - y'| > 2$ we get

$$\begin{aligned} \mathbb{P}(|X_t^{y,0} - X_t^{y',0}| = 0 \text{ for some } t \geq 0) &\leq \mathbb{P}(|X_t^{y,0} - X_t^{y',0}| < 2 \text{ for some } t \geq 0) \\ &\leq \mathbb{P}(|Y_t^{y,0} - Y_t^{y',0}| < 2 \text{ for some } t \geq 0) \\ &\leq \frac{\underline{R}^{d-2}}{|y - y'|^{d-2}}. \end{aligned}$$

Last inequality is obtained from the fact that $|Y_t - Y'_t|^{2-d}$ is a local martingale and standard estimates involving hitting times of balls centered at the origin. Fix $\underline{R} > 2, n \geq 2, \varepsilon > 0$ and $g(n, d, \varepsilon, \underline{R})$, to be defined later. Since for the Poisson process we can always choose a point exterior to any bounded region, choose $x_1, \dots, x_n \in \eta_0$ iteratively from η_0 , a configuration picked from the invariant measure for the coalescing random walks, in the following way: Pick x_1 arbitrary, x_i at distance bigger than g of the $i - 1$ points chosen before, $i = 1, \dots, n$. Then

$$\begin{aligned} & \mathbb{P}(\text{there are less than } n \text{ walkers that do not meet}) \\ & \leq \mathbb{P}(\exists i, j \in \{1, \dots, n\} \text{ with } |X_t^{x_i, 0} - X_t^{x_j, 0}| = 0 \text{ for some } t \geq 0) \\ & \leq \sum_{i, j} \mathbb{P}(|X_t^{x_i, 0} - X_t^{x_j, 0}| = 0 \text{ for some } t \geq 0) \\ & \leq \frac{n(n-1)}{2} \frac{\underline{R}^{d-2}}{|g(n, d, \varepsilon, \underline{R})|^{d-2}} \leq \varepsilon \end{aligned}$$

for a sufficiently large $g(n, d, \varepsilon, \underline{R})$. Hence

$$\mathbb{P}(\text{there are a finite number of walkers that do not meet}) = 0.$$

(d) Let $\xi_t^r \subset \eta_t$ be the set of points of η_t with branches going up to time $r < t$. Obviously ξ_t^r is a decreasing set in r : $\xi_t^r \subset \xi_t^{r'}$ for $r \leq r' \leq t$. Hence the set $\xi_t = \lim_{r \rightarrow -\infty} \xi_t^r$ is well defined and describes the set of points having infinite branches backwards in time. By construction the law of ξ_t does not depend on t and it is stochastically bounded by the law of η_t . The process ξ_t is stationary by construction and Markovian with generator

$$L_o f(\xi) = \lambda \int_{\mathbb{R}^{d-1}} [f(\{\xi \setminus B(u)\} \cup \{u\}) - f(\xi)] \mathbf{1}\{\xi \cap B(u) \neq \emptyset\} du,$$

for functions $f : \mathcal{S}_{d-1} \rightarrow \mathbb{R}$ depending on bounded regions. It describes the motion of coalescing random walks (without births).

We want to show that $\xi_t = \emptyset$ with probability one: $\mathbb{P}(\xi_t = \emptyset) = 1$ for all t in \mathbb{R} .

Assume by contradiction that $\mathbb{P}(|\xi_t| \geq 1) > 0$, where $|\xi|$ stands for the total number of particles in ξ . We claim that $\mathbb{P}(|\xi_t| = 1) + \mathbb{P}(\xi_t = \emptyset) < 1$. Indeed, if $\mathbb{P}(|\xi_t| = 1) + \mathbb{P}(\xi_t = \emptyset) = 1$, for any finite set A , and any $t \in \mathbb{R}, t' \geq 0$,

$$\mathbb{P}(\xi_{t+t'} \cap A = \emptyset) = \mathbb{P}(\xi_t = \emptyset) + \mathbb{P}(\xi_{t+t'} \cap A = \emptyset \mid |\xi_t| = 1) \mathbb{P}(|\xi_t| = 1).$$

If $|\xi_t| = 1$, for $t' \geq 0$, $\xi_{t+t'}$ is a symmetric random walk on \mathbb{R}^{d-1} . In particular, $\mathbb{P}(\xi_{t+t'} \cap A = \emptyset \mid |\xi_t| = 1)$ converges to 1 as $t' \uparrow \infty$. Hence, by stationarity,

$$\mathbb{P}(\xi_t \cap A = \emptyset) = \mathbb{P}(\xi_t = \emptyset) + \mathbb{P}(|\xi_t| = 1) = 1$$

for all sets A . Letting $A \uparrow \mathbb{R}^{d-1}$, we obtain that $\mathbb{P}(\xi_t = \emptyset) = 1$ in contradiction with the first assumption.

We just proved that $\mathbb{P}(|\xi_t| = 1) + \mathbb{P}(\xi_t = \emptyset) < 1$ so that $\mathbb{P}(|\xi_t| \geq 2) > 0$. In particular, there exists a bounded set A such that $\mathbb{P}(|A \cap \xi_t| \geq 2) > 0$.

Let Λ be a $(d - 1)$ -dimensional cubic box centered at the origin. Let $\beta_t = \mathbb{E}(|\xi_t \cap \Lambda|)$. By invariance of the construction of ξ_t , β_t does not depend on t . On the other hand, β_t grows at most at rate

$$\lambda |\delta \Lambda| = \lambda \mathcal{O}(|\Lambda|^{(d-2)/(d-1)})$$

corresponding to the entrance of infinite branches through the boundary of Λ . This happens at rate λ times the volume of the set of points in Λ being at distance less than one from some point outside Λ ; this volume is of the order $|\Lambda|^{(d-2)/(d-1)}$.

On the other hand, β_t decreases at least at rate

$$\lambda \int_{\Lambda} \mathbb{P}(B(u) \cap \xi_t \geq 2) du = \lambda \mathbb{P}(B(0) \cap \xi_t \geq 2) |\Lambda|,$$

where we recall $B(u)$ is the $(d - 1)$ -euclidean ball of radius 1 centered at u . The identity follows from the translation invariance of ξ_t . We showed above that there exists a bounded set A such that $\mathbb{P}(|A \cap \xi_t| \geq 2) > 0$. Since in a finite amount of time one can find Poisson points (x, r) taking the (at least) two points in A at distance less than one, $\mathbb{P}(B(u) \cap \xi_t \geq 2) > 0$. Therefore, β_t decreases at least by a constant times $\lambda|\Lambda|$. We conclude that

$$\frac{d}{dt}\beta_t = O(|\Lambda|^{(d-2)/(d-1)}) - O(|\Lambda|) < 0$$

for sufficiently large Λ . This is in contradiction with the time independence of β_t and shows that $\mathbb{P}(\xi_t \neq \emptyset) = 0$.

(f) For almost every point configuration S , two different points in S have all coordinates different. Hence they are distinctly ordered by the first coordinate. We want to show that every walk will eventually coalesce with a younger walk at coalescence time. Recall that $\tau_n(x, r)$ are the jump times of the walk $X_t^{x,r}$. Let K_n be the intersection of the following events:

- Between τ_n and $\tau_n + 1$, $X_t^{x,r}$ does not jump.
- Between τ_n and $\tau_n + 1$ a new walk is born at distance bigger than one and less than two of $X_t^{x,r}$, call it $X_t^{x',r'}$; $|x' - X_t^{x,r}| \in (1, 2)$, $r' \in [\tau_n, \tau_n + 1)$.
- Between $\tau_n + 1$ and $\tau_n + 2$ both walks coalesce.

τ_n is a hitting time: the event $\{\tau_n \leq t\}$ is \mathcal{F}_t measurable, i.e. it depends only on the points of the past of t . The above event depends on points in a finite region in the future of τ_n . It has positive probability and given that it has occurred, the probability that $x' \leq X_t^{x,r}$ and $x' \geq X_t^{x,r}$ are equal. In the second case $X_t^{x',r'}$ has coalesced with a younger walk. Since τ_n is a Poisson process, we can make infinitely many independent attempts, each with positive probability of success. Hence with probability one each walk will coalesce with a younger walk. \square

Notice that our proof of (f) uses the independence of disjoint regions of the Poisson process and the fact that a success is attained after a geometric number of independent attempts. Another possibility would be to use ergodic theory *à la* Burton and Keane [4]; see also Holroyd and Peres [9].

5. On point-stationarity

In this section we shall briefly consider Problem 1.3 from the introduction.

Let S° be a point process in \mathbb{R}^d with a point at the origin, $0 \in S^\circ$. According to Thorisson [13,14, Chapter 9], S° is *point-stationary* if for any shift-measurable stationary random field $Y = (Y_u : u \in \mathbb{R}^d)$, which is independent of S° , it holds that

$$S^\circ - \pi(S^\circ, Y) = S^\circ \quad \text{in distribution,}$$

where π is such that the associated point-shift is a bijection. It is further shown that S° is point-stationary if and only if, for all $u \in \mathbb{R}^d$,

$$E[f(u + S^\circ - U) | C^\circ] = E[f(S^\circ - U) | C^\circ], \tag{5.1}$$

where C° is the Voronoi cell of the point at the origin and, conditionally on S° , U is uniformly distributed on C° . In particular when $E[|C^\circ|] < \infty$, this means that the reversed Palm version S of S° is stationary,

$$S - u = S, \quad \text{in distribution for } u \in \mathbb{R}^d.$$

We shall now show that if there exist non-randomized bijective point shifts like those constructed in the introduction for the Poisson process in the cases $d = 2$ and 3, then point-stationarity reduces to invariance under those point

shifts. That is, the stationary independent background field Y is not needed to ensure point-stationarity in those cases.

Theorem 5.1. *Let S° be a point process in \mathbb{R}^d . Let $\pi_n, n \in \mathbb{Z}$, be point maps and θ_{π_n} the associated point shifts. Suppose that*

$$\pi_{-n} \theta_{\pi_n} = -\pi_n, \tag{5.2}$$

that the points $X_n = \pi_n(S^\circ), n \in \mathbb{Z}$, are all distinct, and that $S^\circ = \{X_n: n \in \mathbb{Z}\}$ a.s. If

$$S^\circ - X_n = S^\circ, \quad \text{in distribution for } n \in \mathbb{Z}, \tag{5.3}$$

then S° is point-stationary.

Proof. Let C^n be the Voronoi cell of X_n . Due to (5.2) and (5.3),

$$\int_{(u+C^\circ) \cap C_n} f(s + S^\circ) ds = \int_{(u+C^{-n}) \cap C_0} f(s + S^\circ) ds$$

in distribution. Take expectations and sum over n to obtain

$$\mathbb{E} \left[\int_{u+C^\circ} f(s + S^\circ) ds \right] = \mathbb{E} \left[\int_{C^\circ} f(s + S^\circ) ds \right].$$

This is a reformulation of (5.1) which is equivalent to point-stationarity, and the proof is complete. \square

6. Final remarks and open problems

We have deterministically constructed connected trees with a unique infinite succession line having as vertices the points of a homogeneous Poisson process in dimensions $d = 2$ and $d = 3$. This naturally poses the following problems.

- (a) Construct an infinite succession line for a Poisson process in dimension $d \geq 4$. Due to Theorem 5.1, this would in particular solve Problem 1.2.
- (b) Construct an infinite succession line independent of the choice of a direction. The minimal spanning tree construction has this property, but it is still to be proven that it possesses a unique infinite succession line. On the other hand the minimal spanning tree has been proven to be connected only in dimension 2.

Holroyd and Peres [9] have solved (a) and (b) in all dimensions by constructing a connected tree with finite branches having as vertices the points of the Poisson process in a translation and rotationally invariant way.

A challenging problem is to give sufficient conditions for an ergodic translation invariant point process in any dimension to be ordered in a deterministic way in a unique infinite succession line. This would be an important step towards solving Problem 1.3.

The infinite succession line in $d = 2$ corresponds to the “random maze” introduced in the Fig. 2 of Tóth and Werner [15]. Propositions 3.1 and Lemmas 3.2 and 3.4 of that paper together with the convergence of the tree to the Brownian web imply that the infinite succession line converges to a random line that totally order the two-dimensional real numbers.

Acknowledgements

We thank Dana Randall and Xian-Yuan Wu for calling our attention to the minimal spanning tree, Ken Alexander for his comments and references regarding the minimal spanning tree, Balint Tóth for enlightening discussions about succession lines, Anish Sarkar for the references on drainage networks and Yuval Peres for indicating the reference to Häggström and Meester.

This paper is partially supported by FAPESP, CNPq, PRONEX.

References

- [1] K.S. Alexander, Percolation and minimal spanning forests in infinite graphs, *Ann. Probab.* 23 (1) (1995) 87–104.
- [2] R. Arratia, Coalescing Brownian motions and the voter model on \mathbb{Z} , Unpublished manuscript, 1981, available from rarratia@math.usc.edu.
- [3] R. Arratia, Coalescing Brownian motions on the line, Phd Thesis. Univ. of Madison, Wisconsin, 1981.
- [4] R.M. Burton, M.S. Keane, Density and uniqueness in percolation, *Comm. Math. Phys.* 121 (1989) 501–505.
- [5] P.A. Ferrari, L.R.G. Fontes, X.-Y. Wu, Poisson trees converge to Brownian web, <http://arxiv.org/abs/math.PR/0304247>, 2003.
- [6] L.R.G. Fontes, M. Isopi, C.M. Newman, K. Ravishankar, The Brownian web: characterization and convergence, <http://arxiv.org/abs/math.PR/0304119>, 2003.
- [7] S. Gangopadhyay, R. Roy, A. Sarkar, Random oriented trees: a model of drainage networks, Preprint Indian Statistical Institute isid/ms/2002/05, <http://www.isid.ac.in/statmath/eprints/2002/isid200205.pdf>, 2002.
- [8] O. Häggström, R. Meester, Nearest neighbor and hard sphere models in continuum percolation, *Random Structures Algorithms* 9 (3) (1996) 295–315.
- [9] A.E. Holroyd, Y. Peres, Trees and matchings from point processes, <http://arxiv.org/abs/math.PR/0211455>, 2003.
- [10] C.M. Newman, D.L. Stein, Multiple states and thermodynamic limits in short-ranged Ising spin-glass models, *Phys. Rev. B* 46 (1992) 973–982.
- [11] C.M. Newman, D.L. Stein, Spin-glass model with dimension-dependent ground state multiplicity, *Phys. Rev. Lett.* 72 (1994) 2286–2289.
- [12] I. Rodriguez-Iturbe, A. Rinaldo, *Fractal River Networks: Chance and Self-Organization*, Cambridge University Press, New York, 1997.
- [13] H. Thorisson, Point-stationarity in d dimensions and Palm theory, *Bernoulli* 5 (5) (1999) 797–831.
- [14] H. Thorisson, *Coupling, Stationarity, and Regeneration. Probability and its Applications*, Springer-Verlag, New York, 2000.
- [15] B. Tóth, W. Werner, The true self-repelling motion, *Probab. Theory Related Fields* 111 (3) (1998) 375–452.