Adv. Appl. Prob. **50,** 1275–1293 (2018)
*doi:*10.1017/apr.2018.60
© *Applied Probability Trust* 2018

NEAREST-NEIGHBOUR MARKOV POINT PROCESSES ON GRAPHS WITH EUCLIDEAN EDGES

M. N. M. VAN LIESHOUT,* *CWI and University of Twente*

Abstract

We define nearest-neighbour point processes on graphs with Euclidean edges and linear networks. They can be seen as analogues of renewal processes on the real line. We show that the Delaunay neighbourhood relation on a tree satisfies the Baddeley–Møller consistency conditions and provide a characterisation of Markov functions with respect to this relation. We show that a modified relation defined in terms of the local geometry of the graph satisfies the consistency conditions for all graphs with Euclidean edges that do not contain triangles.

Keywords: Delaunay neighbour; graph with Euclidean edges; linear network; Markov point process; nearest-neighbour interaction; renewal process

2010 Mathematics Subject Classification: Primary 60G55
Secondary 60D05

1. Introduction

In recent years a theory of point processes on linear networks has been emerging so as to be able to analyse, for example, the prevalence of accidents on motorways, the occurrence of street crimes, and other data described in the first chapter of the pioneering monograph by Okabe and Sugihara [12]. Although there exists a mature theoretical framework for point processes on Euclidean spaces [7], [8], the development of a similar theory on linear networks is complicated by the geometry inherent in the network. In particular, it is not possible to define strictly stationary models, as the network may not be closed under translations. For this reason, most attention has focussed on the development of second-order summary statistics [13].

Little attention has been paid to model building, with a few notable exceptions. The first serious work in this direction seems to be that by Baddeley *et al.* [5], who constructed certain types of Cox processes as well as a Switzer type and a cell process. The authors concluded that familiar procedures for constructing models tend not to produce processes on a linear network that are pseudostationary with respect to the shortest path distance, except when the network is a tree—an unrealistic assumption for a road network. Another important contribution is the work by Anderes *et al.* [1] who expanded the modelling framework in various directions. They relaxed the assumption of [5], [12], and [13] that a linear network consists of a finite union of straight line segments which intersect only at the vertices, in the sense that the segments are replaced by parametrised rectifiable curves that may or may not overlap. The parametrisations have the additional advantage of naturally defining a weighted shortest path distance. In the motivating example where the linear network represents a road network, such a generalisation allows for bridges or tunnels and for distance to be measured in travel time where appropriate.

Received 24 October 2017; revision received 10 August 2018.
* Postal address: CWI, PO Box 94079, NL-1090 GB Amsterdam, The Netherlands.
Email address: marie-colette.van.lieshout@cwi.nl

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

Additionally, Anderes *et al.* [1] constructed log-Gaussian Cox processes in terms of a Gaussian process on the network specified by an isotropic covariance function. In the expanded paper [2], special attention was paid to alternative metrics.

The Cox models discussed above are clustered in nature, that is, exhibit a positive association between the points. Moreover, their intensity function and pair correlation function are explicitly known, and pseudostationarity is defined to mean that the intensity function is constant, and that the pair correlation function of a pair of points depends only on the distance between the two points. In this paper our aim is to develop appropriate analogues of renewal processes, exploiting the one-dimensional nature of a network. Recall that renewal processes exhibit the property that the probability (in an infinitesimal sense) of an event at a given location conditional on the realisation of the process elsewhere depends only on the two nearest points, regardless of how far away they may be. In general spaces such models are known as nearest-neighbour Markov point processes [3]. For the related class of fixed-range Markov point processes introduced by Ripley and Kelly [14], the conditional probabilities of finding a point at a given location depends on the configuration only through points within a prescribed distance from that location. Note that Markovian models are particularly, but not exclusively, suited to model inhibition [15].

In contrast to Cox models, the second-order summary statistics of Markov and nearest-neighbour Markov point processes may not be available in closed form. Instead, explicit expressions exist for the conditional intensities and the likelihood can often be expressed as a product of interaction functions. Thus, an appropriate concept of pseudostationarity in this context is that the interaction functions depend on the interpoint distances only. We refer the reader to [11] for an overview and critique of techniques to define inhomogeneity in point process models.

The plan of this paper is as follows. In Section 2 we recall the definitions of Anderes *et al.* [1], [2] regarding graphs with Euclidean edges, the weighted shortest path metric thereon, and the Poisson process defined on them. We go on to define point processes in terms of their probability density with respect to a Poisson process and note that the theory of fixed-range Markov point processes carries over immediately. In Section 3 we extend the notion of a Markov point process with respect to the Delaunay nearest-neighbour relation [3] to graphs with Euclidean edges and state our main results. More specifically, we show that the Delaunay relation on a tree satisfies the Baddeley–Møller consistency conditions and we provide a characterisation of the Markov functions with respect to this relation. We then use the graph structure to define a modified Delaunay relation and show that it satisfies the consistency conditions on a wide class of graphs with Euclidean edges. The proofs are given in Sections 4 and 5.

2. Preliminaries

2.1. Graphs with Euclidean edges

In their pioneering monograph on the subject, Okabe and Sugihara [12, p. 31] defined a network as a finite union

$$L = \bigcup_{i=1}^{n} l_i, \qquad n \in \mathbb{N},$$

of straight line segments l_i in $\mathbb{R}^2$ or $\mathbb{R}^3$ that intersect only at their endpoints in such a way that L is connected. The representation is not unique since a line segment may arbitrarily be split into two pieces without affecting the union L.

A more general definition is given by Anderes *et al.* [1]. They replaced the straight line segments by curves parametrised by bijections. We impose the further condition that these parametrisations are continuously differentiable.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

Definition 2.1. A graph with Euclidean edges in $\mathbb{R}^2$ is a triple $G = (V, E = (e_i)_i, \Phi = (\phi_i)_i)$ such that

- (V, E) is a finite, simple connected graph, that is, has neither loops nor multiple edges;
- every edge $e_i = \{v_i^1, v_i^2\} \in E$, $v_i^1, v_i^2 \in V$, is parametrised by the inverse of homeomorphism ϕ_i, and $\phi_i^{-1}\colon J_i \to \mathbb{R}^2$ for an open interval $\varnothing \neq J_i \subset \mathbb{R}$ with endpoints $\phi_i(v_i^j)$, $j = 1, 2$, is continuously differentiable.

In other words, an edge is associated with a set $\phi_i^{-1}(J_i) \subset \mathbb{R}^2$ that does not contain the endpoints. The continuity of ϕ_i^{-1} implies that the set is connected.

A graph with Euclidean edges gives rise to the space

$$L = (\{0\} \times V) \cup \bigcup_{i=1}^{n(E)} (\{i\} \times \phi_i^{-1}(J_i)),$$

where $n(E) < \infty$ is the cardinality of E. The labels i serve to identify the edges and will prevent paths from 'jumping from one edge to another' in case their interiors overlap. For instance, if L represents a road network, overlap is typically present due to tunnels and bridges [1].

As an aside, if there is no overlap between the edge interiors, one may drop the labels and simply consider the disjoint union

$$V \cup \bigcup_{i=1}^{n(E)} \phi_i^{-1}(J_i),$$

which in turn reduces to the classic linear networks if all edges $\phi_i^{-1}(J_i)$ are straight line segments. From now on, we will work with the general space L including the labels.

2.2. Weighted shortest path metric

The family Φ of parametrisations that is part of the definition of a graph with Euclidean edges can be used to define concepts of length and distance [1].

Definition 2.2. Let (V, E, Φ) be a graph with Euclidean edges. For every $i = 1, \dots, n(E)$, define the ϕ_i^{-1}-induced length measure on the σ-algebra $\{\phi_i^{-1}(B)\colon B \subset \bar{J}_i, B \text{ Borel}\}$ on $\phi_i^{-1}(\bar{J}_i)$ generated by the functions ϕ_i as follows. The induced length of the set $\phi_i^{-1}(B)$ is equal to the Euclidean length of B in the closed interval $\bar{J}_i$. In particular, the edge $e_i = \{v_i^1, v_i^2\}$ has length $|\phi_i(v_i^1) - \phi_i(v_i^2)|$ equal to the Euclidean length of $\bar{J}_i$. The Φ-induced length is the ensemble of ϕ_i^{-1}-induced length measures.

Since the parametrisations ϕ_i^{-1} are continuously differentiable, the edges they define have finite arc lengths. Indeed, under the arc length parametrisation, the induced length in the sense of Definition 2.2 corresponds to the arc length [5]. A different parametrisation may induce a different length measure.

Definition 2.3. Let (V, E, Φ) be a graph with Euclidean edges. A walk between two elements (i, x) and (j, y) of L travels alternatingly from (i, x) to (j, y) along a finite number of vertices and edges such that the vertices are endpoints of the edges and edge labels do not change in between vertices. If $i \neq 0$, the first part of the walk travels from x along $\phi_i^{-1}(J_i)$; similarly, for $j \neq 0$, the last part of the walk is along $\phi_j^{-1}(J_j)$ to y. In particular, a walk part between two points along the same curve $\phi_i^{-1}(J_i)$ does not reverse its tracks.

Definition 2.4. Let (V, E, Φ) be a graph with Euclidean edges. A path between two different elements (i, x) and (j, y) of L is a walk in which all edge segments and all vertices are different.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

The weight of the path is the sum of the Φ-induced lengths of the edge segments in it, and the weighted shortest path distance $d_G((i,x),(j,y))$ between two elements (i,x) and (j,y) of L is the smallest weight carried by any path between them.

The weighted shortest path distance on a graph with Euclidean edges defines a metric. Note that it is possible that the shortest weighted path between two vertices joined by an edge is not along that edge! For further details, see [1].

2.3. Poisson processes on graphs with Euclidean edges

Our goal in this paper is to define point processes on graphs with Euclidean edges in terms of a density with respect to a unit-rate Poisson process [15]. To this end, we recall the definition of a Poisson process on a graph with Euclidean edges given by Anderes *et al.* [1].

Definition 2.5. Let $G = (V, E = (e_i)_i, \Phi = (\phi_i)_i)$ be a graph with Euclidean edges and let L be the corresponding network. The Lebesgue–Stieltjes measure λ_G is defined as follows. For every i and every set B_i in the σ-algebra on $\phi_i^{-1}(J_i)$ generated by ϕ_i, set

$$\lambda_G^i(B_i) = \int_{J_i} \mathbf{1}\{\phi_i^{-1}(t) \in B_i\} \left| \frac{\mathrm{d}}{\mathrm{d}t} \phi_i^{-1}(t) \right| \mathrm{d}t,$$

where $|\cdot|$ denotes the norm of the gradient. Then

$$\lambda_G\left(\bigcup_{i=1}^{n(E)} (\{i\} \times B_i) \right) = \sum_{i=1}^{n(E)} \lambda_G^i(B_i)$$

is a measure on L equipped with the finite disjoint union σ-algebra in which a set is measurable if and only if it can be written as $\bigcup_{i=0}^{n(E)} (\{i\} \times B_i)$ with B_i in the σ-algebra generated by ϕ_i for $i > 0$ and the power set of V for $i = 0$. By default, the set $\{0\} \times V$ has Lebesgue–Stieltjes measure zero.

Note that Definition 2.5 does not depend on the choice of parametrisations.

Definition 2.6. Let $G = (V, E = (e_i)_i, \Phi = (\phi_i)_i)$ be a graph with Euclidean edges and let L be the corresponding network. The unit-rate Poisson process on L is defined as follows. For every i and every set B_i in the σ-algebra on $\phi_i^{-1}(J_i)$ generated by ϕ_i, independently for different i,

- the number of points in $\{i\} \times B_i$ is Poisson distributed with expectation $\lambda_G(\{i\} \times B_i)$;
- given that the number of points falling in $\{i\} \times B_i$ is n_i, these n_i points are independent and identically distributed according to the probability density function $1/\lambda_G(\{i\} \times B_i)$.

In words, the unit-rate Poisson process scatters a Poisson number of points independently and uniformly on every edge, and the average number of such points is equal to the arc length of the edge. Note that this point process is simple, that is, almost surely, there are no multiple points. Moreover, the definition does not depend on the choice of parametrisations.

The integral of a measurable function $f : \mathbb{R}^2 \to \mathbb{R}^+$ with respect to the Lebesgue–Stieltjes measure λ_G is defined as the sum of the line integrals of f along the rectifiable curves parametrised by the functions $\phi_i^{-1} : J_i \to \mathbb{R}^2$:

$$\int_L f \, \mathrm{d}\lambda_G = \sum_{i=1}^{n(E)} \int_{J_i} f(\phi_i^{-1}(t)) \left| \frac{\mathrm{d}}{\mathrm{d}t} \phi_i^{-1}(t) \right| \mathrm{d}t.$$

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

Higher-order integrals are defined analogously. Note that the definition does not depend on the choice of parametrisations.

Example 2.1. Consider the functions f_G^i defined on $\phi_i^{-1}(J_i)$ by

$$f_G^i(x) = \left| \left(\frac{\mathrm{d}}{\mathrm{d}t} \phi_i^{-1} \right) (\phi_i(x)) \right|^{-1}.$$

Then, for B_i in the σ-algebra generated by ϕ_i,

$$\begin{aligned} \int_L \mathbf{1}\{x \in B_i\} f_G^i(x) \, \mathrm{d}\lambda_G(x) &= \int_{J_i} \mathbf{1}\{\phi_i^{-1}(t) \in B_i\} \left| \left(\frac{\mathrm{d}}{\mathrm{d}t} \phi_i^{-1} \right) (\phi_i(\phi_i^{-1}(t))) \right|^{-1} \left| \frac{\mathrm{d}}{\mathrm{d}t} \phi_i^{-1}(t) \right| \mathrm{d}t \\ &= |\phi_i(B_i)|, \end{aligned}$$

the Euclidean length of $\phi_i(B_i)$. In other words, the functions f_G^i, $i = 1, \ldots, n(E)$, define the weighted shortest path distance d_G on L.

A simple point process X on L is said to have probability density p with respect to the unit-rate Poisson process if $\mathbb{P}(X \in F)$ is equal to

$$\sum_{n=0}^{\infty} \frac{\mathrm{e}^{-\lambda_G(L)}}{n!} \int_L \cdots \int_L \mathbf{1}\{\{x_1, \ldots, x_n\} \in F\} p(\{x_1, \ldots, x_n\}) \, \mathrm{d}\lambda_G(x_1) \cdots \mathrm{d}\lambda_G(x_n) \tag{2.1}$$

for all F in the usual σ-algebra on finite-point configurations in L generated by the counts $X(B)$ with B in the finite disjoint union σ-algebra on L [7], [8]. If the probability density p is such that the ratio

$$\frac{p(\{x_1, \ldots, x_n, x_{n+1}\})}{p(\{x_1, \ldots, x_n\})},$$

when well defined, depends only on those x_j, $j = 1, \ldots, n$, for which $d_G(x_{n+1}, x_j) \le r$ for some fixed $r > 0$, X is fixed-range Markov and the Hammersley–Clifford theorem [14] applies. For the sake of convenience, we will use the notation $\boldsymbol{x} = \{x_1, \ldots, x_n\}$ for configurations of finitely many distinct points in L.

For further details on graphs with Euclidean edges, we refer the reader to [1] and [2].

3. Nearest-neighbour Markov point processes

The purpose of this section is to define appropriate analogues of the well-known class of renewal processes on the real line. Intuitively speaking, we are looking for a class of point processes in which the conditional behaviour at a given point depends on the remainder of the configuration only through its 'nearest neighbours'. We shall show that in the case that the network L is a tree, the weighted shortest path distance may be used to define which points are each other's nearest neighbours. In the general case, we will use the geometry of the network to define a local neighbourhood relation.

3.1. The Delaunay relation

In this section we adapt Baddeley and Møller's definition of configuration-dependent neighbourhoods and cliques in Euclidean spaces [3] to our context.

Definition 3.1. Let (V, E, Φ) be a graph with Euclidean edges and let L be the corresponding network. Let $\boldsymbol{x} \subset L$ be a finite configuration of distinct points, and define the Delaunay relation '$\sim_{\boldsymbol{x}}$' as the symmetric, reflexive relation on $\boldsymbol{x}$ given by

$$x_1 \sim_{\boldsymbol{x}} x_2 \quad \Longleftrightarrow \quad C(x_1 \mid \boldsymbol{x}) \cap C(x_2 \mid \boldsymbol{x}) \neq \varnothing,$$

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

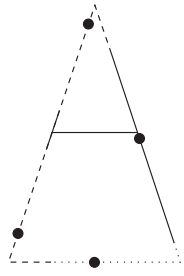


FIGURE 1: Points on a linear network. The Voronoi cells are indicated by solid, dashed, and dotted lines.

where, for $x_j \in \boldsymbol{x}$, $j = 1, 2$,

$$C(x_j \mid \boldsymbol{x}) = \{(i, y) \in L : d_G((i, y), x_j) \leq d_G((i, y), x) \text{ for all } x \in \boldsymbol{x}\}$$

is the Voronoi cell of x_j in $\boldsymbol{x}$. The $\boldsymbol{x}$–neighbourhood of a subset $\boldsymbol{z} \subseteq \boldsymbol{x}$ is defined as

$$N(\boldsymbol{z} \mid \boldsymbol{x}) = \{x \in \boldsymbol{x} : x \sim_{\boldsymbol{x}} z \text{ for some } z \in \boldsymbol{z}\}.$$

The configuration $\boldsymbol{z}$ is an $\boldsymbol{x}$-clique if, for each $z_1, z_2 \in \boldsymbol{z}$, $z_1 \sim_{\boldsymbol{x}} z_2$. By convention, the empty set and singletons are cliques too.

We shall use the notation $y \in L$ and $(i, y) \in L$ as the occasion demands, that is, we will refrain from mentioning i explicitly unless it is necessary to do so.

To illustrate the definitions, consider Figure 1. It depicts four points, shown as filled circles, on a linear network. The corresponding four Voronoi cells are indicated by solid, dashed, and dotted lines. Due to the cycle, three different line types are needed. Note that the Voronoi cells may extend over various line segments and branch off at vertices (e.g. for the cell indicated by a solid line). The number of Delaunay neighbours varies per point. The topmost point has a single neighbour, the rightmost point has three, and the other points each have two neighbours. As for the clique sizes, the empty set, singletons, and pairs of points whose Voronoi cells intersect are cliques. Additionally, the three lower points that lie on a cycle in the network form a clique of size three. There is no clique of cardinality 4.

Some elementary properties of the relation are collected in the next lemma.

Lemma 3.1. *Let (V, E, Φ) be a graph with Euclidean edges and let L be the corresponding network. The Delaunay relation on L satisfies the following properties for all finite configurations of distinct points $\boldsymbol{x} \subset L$:*

- *if $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 1$ then also $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$ for all $\boldsymbol{y} \subseteq \boldsymbol{z} \subseteq \boldsymbol{x}$;*
- *if $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 0$ then also $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 0$ for all $\boldsymbol{y} \subseteq \boldsymbol{z} \subseteq \boldsymbol{x}$.*

Here $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 1$ if $\boldsymbol{y}$ is an $\boldsymbol{x}$-clique and 0 otherwise.

Proof. See Section 4. □

3.2. The Delaunay relation on trees

Recall that (V, E) is said to be a tree if it has no cycles, that is, there is no closed path $(v_0, v_1, \ldots, v_p, v_0)$, $v_i \in V$, of positive length $(p > 0)$. A graph with Euclidean vertices $G = (V, E, \Phi)$ is a tree if (V, E) is a tree.

It is well known that a graph is a tree if and only if there is exactly one path between any two vertices [6], a property that is inherited by the network associated with G.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

Lemma 3.2. *A graph with Euclidean edges is a tree if and only if there is exactly one path between any two points (i, x) and (j, y) in L.*

Proof. See Section 4. □

We are particularly interested in the geometrical arrangement of the paths between three points.

Lemma 3.3. *Let (V, E, Φ) be a graph with Euclidean edges that is a tree and let L be its associated network. Consider a triple $\boldsymbol{y} = \{y_1, y_2, y_3\} \subset L$ of distinct points. Then there exist unique paths between the elements of $\boldsymbol{y}$ that*

- *either form a three-pointed star with sides of strictly positive length emanating from a vertex $(0, v) \in L$; or*
- *combine into a single path.*

Proof. See Section 4. □

A finite configuration $\boldsymbol{y} \subset L$ is said to be in general position if no three points lie on the boundary of the same d_G-ball. Clearly, the class of all finite configurations in general position is hereditary.

Lemma 3.4. *Let (V, E, Φ) be a graph with Euclidean edges that is a tree and let L be its associated network. Then the clique sizes with respect to the Delaunay relation are at most two on the class of finite configurations in general position. Moreover, for all $\boldsymbol{y} = \{y_1, y_2\} \subseteq \boldsymbol{x}$ with $\boldsymbol{x}$ in general position, $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 1$ if and only if the midpoint of $\boldsymbol{y}$ with respect to the weighted shortest path metric along the unique path between y_1 and y_2 lies in $C(y_1 \mid \boldsymbol{x}) \cap C(y_2 \mid \boldsymbol{x})$.*

Proof. See Section 4. □

In order to define Markov functions, consistency conditions must be imposed on the family of neighbourhood relations [3].

Definition 3.2. Let (V, E, Φ) be a graph with Euclidean edges and let L be its associated network. Consider finite configurations $\boldsymbol{y} \subseteq \boldsymbol{z} \subset L$ of distinct points and points $u, v \in L$ such that $u, v \notin \boldsymbol{z}$. Then the Baddeley–Møller consistency conditions read as follows:

(C1) $\chi(\boldsymbol{y} \mid \boldsymbol{z}) \neq \chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u\})$ implies that $\boldsymbol{y} \subseteq N(\{u\} \mid \boldsymbol{z} \cup \{u\})$;

(C2) if $u_1 \not\sim_{\boldsymbol{x}} u_2$ for $\boldsymbol{x} = \boldsymbol{z} \cup \{u_1, u_2\}$ then

$$\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1\}) + \chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_2\}) = \chi(\boldsymbol{y} \mid \boldsymbol{z}) + \chi(\boldsymbol{y} \mid \boldsymbol{x}).$$

To see the rationale behind these conditions, suppose that we wish to build a pairwise interaction model by setting

$$p(\boldsymbol{x}) = \alpha \prod_{x \sim_{\boldsymbol{x}} y} \gamma(x, y)$$

for some Borel measurable function $\gamma(\cdot, \cdot) \geq 0$ that is symmetric in its arguments. Then, for such functions to be Markovian, the ratio

$$\frac{p(\boldsymbol{x} \cup \{u\})}{p(\boldsymbol{x})} = \prod_{x \sim_{\boldsymbol{x} \cup \{u\}} u} \gamma(x, u) \frac{\prod_{x \sim_{\boldsymbol{x} \cup \{u\}} y} \gamma(x, y)}{\prod_{x \sim_{\boldsymbol{x}} y} \gamma(x, y)}$$

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

should depend only on the neighbours of u in '$\sim_{\boldsymbol{x}\cup\{u\}}$' (for a precise formulation, see Definition 3.3 below). Hence, any pair $\{x, y\}$ for which $\chi(\{x, y\} \mid \boldsymbol{x}) \neq \chi(\{x, y\} \mid \boldsymbol{x} \cup \{u\})$ should consist of neighbours of u in $\boldsymbol{x} \cup \{u\}$, as required by (C1). Condition (C2) is needed to ensure that, reversely, every Markov function can be factorised as a product of interaction functions (cf. Theorem 3.2).

In general, the consistency conditions of Definition 3.2 do not hold, as illustrated by the following counterexample.

Example 3.1. Consider the linear network depicted in Figure 1. Let z be the clique consisting of the three lower points in the network. Then $\chi(z \mid z) = 1$. However, placing any additional point $u_1 \notin z$ on the cycle that contains z but not the fourth point in the network would split up two of the points in z. Hence, $\chi(z \mid z \cup \{u_1\}) = 0$ even though the third point of z is no Delaunay neighbour of u_1 in $z \cup \{u_1\}$. Consequently, (C1) does not hold. Upon the addition of a second point u_2 to the cycle in such a way that it is not a Delaunay neighbour of u_1 in the resulting five-point configuration, the identities $\chi(z \mid z \cup \{u_1, u_2\}) = \chi(z \mid z \cup \{u_1\}) = \chi(z \mid z \cup \{u_2\}) = 0$ and $\chi(z \mid z) = 1$ hold, in violation of (C2).

The main theorem of this section is the following.

Theorem 3.1. *Let (V, E, Φ) be a graph with Euclidean edges that is a tree and let L be its associated network. Then the Delaunay relation satisfies (C1) and (C2) on the family of finite configurations in general position.*

Proof. See Section 5. □

3.3. Markov functions

We are now ready to define Markov functions on graphs with Euclidean edges, in analogy with the spatial models of [3].

Definition 3.3. Let (V, E, Φ) be a graph with Euclidean edges and let L be the corresponding network. Let '$\sim_{\boldsymbol{x}}$' be a family of reflexive, symmetric relations on finite configurations $\boldsymbol{x}$ of distinct points in L. Then a function p from the set of finite configurations in general position into $[0, \infty)$ is a *Markov function* with respect to '$\sim_{\boldsymbol{x}}$' if, for all $\boldsymbol{x}$ in general position such that $p(\boldsymbol{x}) > 0$,

(a) $p(\boldsymbol{y}) > 0$ for all $\boldsymbol{y} \subseteq \boldsymbol{x}$;

(b) for all $u \in L \setminus \boldsymbol{x}$ such that $\boldsymbol{x} \cup \{u\}$ is in general position, $p(\boldsymbol{x} \cup \{u\})/p(\boldsymbol{x})$ depends only on u, on $N(\{u\} \mid \boldsymbol{x} \cup \{u\}) \cap \boldsymbol{x} = \{x \in \boldsymbol{x} : x \sim_{\boldsymbol{x}\cup\{u\}} u\}$, and on the relations '$\sim_{\boldsymbol{x}}$' and '$\sim_{\boldsymbol{x}\cup\{u\}}$' restricted to $N(\{u\} \mid \boldsymbol{x} \cup \{u\}) \cap \boldsymbol{x}$.

If (a) holds, p is said to be *hereditary*. The next theorem provides a Hammersley–Clifford factorisation. Similar results for spatial point processes in Euclidean spaces can be found in [3], [4], [10], and [14]. Recall that a function γ from the space of finite-point configurations in general position into $[0, \infty)$ gives rise to an interaction function $\Phi(\boldsymbol{y} \mid \boldsymbol{x}) = \gamma(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$ [3] if the following properties hold. If $\gamma(\boldsymbol{x}) > 0$,

(i) $\gamma(\boldsymbol{y}) > 0$ for all $\boldsymbol{y} \subseteq \boldsymbol{x}$, and

(ii) if, additionally, $\gamma(N(\{u\} \mid \boldsymbol{x} \cup \{u\})) > 0$ then $\gamma(\boldsymbol{x} \cup \{u\}) > 0$.

Theorem 3.2. *Let (V, E, Φ) be a graph with Euclidean edges such that (V, E) is a tree and let L be the corresponding network. Let p be a measurable function from the set of finite*

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

configurations in general position into $[0,\infty)$. *Then* p *is a Markov function with respect to the Delaunay relation if and only if*

$$p(\boldsymbol{x}) \propto \begin{cases} \displaystyle\prod_{x_i\in\boldsymbol{x}} \gamma(\{x_i\}) \prod_{\{i<j\,:\,x_i\sim_{\boldsymbol{x}} x_j\}} \gamma(\{x_i,x_j\}) & \text{if } \gamma(\{x_i,x_j\})>0 \text{ for all distinct } x_i,x_j\in\boldsymbol{x},\\ 0 & \text{otherwise} \end{cases}$$

for some measurable, nonnegative function γ, *extended to configurations* $\boldsymbol{x}$ *consisting of three or more points by assigning the value* 1 *whenever* $\gamma(\boldsymbol{y})>0$ *for all subsets* $\boldsymbol{y}\subset\boldsymbol{x}$ *of cardinality at most* 2 *and* 0 *otherwise, that satisfies (i) and (ii).*

Proof. See Section 5. □

If p can be normalised into a probability density, e.g. by assuming that $\gamma(\{x_i\})$ is bounded and $\gamma(\{x_i,x_j\})\le 1$, p is a Markov density and we can define a nearest-neighbour Markov point process by (2.1).

The characterisation allows some flexibility in the choice of γ.

Example 3.2. As we set out to define analogues of renewal processes on graphs, let us first consider the special case $G=(V=\{(a,0),(b,0)\},\ E=(\{(a,0),(b,0)\}),\ \Phi=(\phi_1))$ for $a<b\in\mathbb{R}$ with $\phi_1^{-1}\colon(a,b)\to\mathbb{R}^2$, $\phi_1^{-1}(t)=(t,0)$, embedding (a,b) into the plane. Since there is only a single edge, we may ignore edge labels and simply consider the linear network

$$L=V\cup((a,b)\times\{0\})=[a,b]\times\{0\}.$$

The weighted shortest path metric d_G on L then corresponds to the Euclidean distance between the first coordinates.

Let π be a probability density function on $\mathbb{R}^+$ and write F_π for the corresponding cumulative distribution function. Suppose that $\pi(x)>0$ for all $x\in[0,b-a]$ and $F_\pi(b-a)<1$. Set

$$p(\boldsymbol{x})=\mathrm{e}^{(b-a)}\pi(\min\boldsymbol{x}-a)(1-F_\pi(b-\max\boldsymbol{x}))\prod_{x_i\sim_{\boldsymbol{x}}x_j}\pi(d_G(x_i^1,x_j^1)),$$

where $\min\boldsymbol{x}=\min\{x_i^1\colon x_i=(x_i^1,0)\in\boldsymbol{x}\}$ is the minimal first coordinate in $\boldsymbol{x}$ and, similarly, $\max\boldsymbol{x}=\max\{x_i^1\colon x_i=(x_i^1,0)\in\boldsymbol{x}\}$. By default, $\min\varnothing=a$ and $\max\varnothing=b$. Then p is a Markov function with respect to the Delaunay relation on L with interaction functions $\gamma(\varnothing)=(1-F_\pi(b-a))\exp(b-a)$,

$$\gamma((x,0))=\frac{\pi(x-a)(1-F_\pi(b-x))}{1-F_\pi(b-a)},$$

and

$$\gamma((x,0),(y,0))=\frac{\pi(d_G((x,0),(y,0)))(1-F_\pi(b-a))}{\pi(\max\{x,y\}-a)(1-F_\pi(b-\min\{x,y\}))},$$

for $x,y\in(a,b)$ and $\gamma(\boldsymbol{z})\equiv 1$ whenever the cardinality of $\boldsymbol{z}$ is bigger than 2. Note that γ is strictly positive. For exponential interarrival densities $\pi(x)=\lambda\exp(-\lambda x)$, the interaction functions $\gamma((x,0))\equiv\lambda$ and $\gamma((x,0),(y,0))\equiv 1$ are constant in accordance with the lack of interaction between points in a Poisson process. In general, the interaction functions may depend on x and y. Moreover, the second-order interaction function is not necessarily a function of $d_G((x,0),(y,0))$ only due to edge effects captured by the denominator.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

Example 3.3. Set $\gamma(\varnothing) = \alpha > 0$ and $\gamma(\{x_i\}) \equiv \beta > 0$, and suppose that, for pairs, γ depends only on the weighted shortest path distance between the members of the pair, that is,

$$\gamma(\{x_i, x_j\}) = g(d_G(x_i, x_j))$$

for some function $g\colon [0, \infty) \to [0, \infty)$. For configurations $\boldsymbol{x}$ with $n(\boldsymbol{x}) > 2$, set

$$\gamma(\boldsymbol{x}) = \mathbf{1}\{g(d_G(x_i, x_j)) > 0 \text{ for all } x_i \neq x_j \in \boldsymbol{x}\}.$$

The function γ thus defined gives rise to an interaction function $\Phi(\boldsymbol{y} \mid \boldsymbol{x}) = \gamma(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$ provided that, for all configurations $\boldsymbol{x}$ and all $u \in L \setminus \boldsymbol{x}$, if the conditions

- $g(d_G(x_i, x_j)) > 0$ for all distinct $x_i, x_j \in \boldsymbol{x}$, and
- $g(d_G(x_i, u)) > 0$ for all $x_i \in \boldsymbol{x}$ such that $x_i \sim_{\boldsymbol{x} \cup \{u\}} u$

hold, then also $g(d_G(x_i, u)) > 0$ for all $x_i \in \boldsymbol{x}$. A sufficient condition is that the function g takes strictly positive values, in which case, using Lemma 3.1,

$$\frac{p(\boldsymbol{x} \cup \{u\})}{p(\boldsymbol{x})} = \beta \frac{\prod_{x_i \sim_{\boldsymbol{x} \cup \{u\}} u} g(d_G(x_i, u))}{\prod_{\{i<j\colon x_i \sim_{\boldsymbol{x}} x_j,\, x_i \not\sim_{\boldsymbol{x} \cup \{u\}} x_j\}} g(d_G(x_i, x_j))}.$$

Example 3.4. Let (V, E, Φ) be a graph with Euclidean edges such that (V, E) is a tree and let L be the corresponding network. Further to Example 3.3, consider the Strauss process defined by $g \equiv \gamma$ for some $\gamma \geq 0$.

When γ and, therefore, g is strictly positive, the unnormalised Strauss density is Markov with respect to the Delaunay relation. Now, a point $(i, x) \in L$ on some edge $\{i\} \times \phi_i^{-1}(J_i)$ has at most two Delaunay neighbours on the same edge and at most one on any other edge; a vertex is related to at most one point on each edge. Therefore, the total number of neighbours of any point on the network is bounded by $n(E) + 1$ and the Strauss density is integrable for all $\gamma > 0$. For $\gamma = 1$, it coincides with the Poisson process upon normalisation.

For $g \equiv 0$, the function $\gamma(\varnothing) = \alpha > 0$, $\gamma(\{x\}) = \beta > 0$, and $\gamma(\boldsymbol{x}) = 0$ for all $\boldsymbol{x}$ with $n(\boldsymbol{x}) \geq 2$ defines an interaction function since the conditions in Example 3.3 are void. The resulting density

$$p(\boldsymbol{x}) = \begin{cases} \alpha\beta^{n(\boldsymbol{x})} & \text{if } n(\boldsymbol{x}) \leq 1, \\ 0 & \text{otherwise} \end{cases}$$

is Markov with respect to the Delaunay relation and may be normalised into a probability density that places mass

$$\frac{1}{1 + \beta\lambda_G(L)}$$

on the empty set and spreads the remaining mass evenly over singletons.

Example 3.5. Let (V, E, Φ) be a graph with Euclidean edges such that (V, E) is a tree and let L be the corresponding network. Further to Example 3.3, consider the hard-core process defined by the function

$$g(r) = \begin{cases} 0 & \text{if } r \leq R, \\ \gamma & \text{otherwise} \end{cases}$$

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

for some fixed $R > 0$ and $\gamma > 0$. For this choice of g, $\gamma(\{x_1, x_2\}) = \gamma$ if $d_G(x_1, x_2) > R$, 0 otherwise, and

$$\gamma(\boldsymbol{x}) = \mathbf{1}\{d_G(x_i, x_j) > R \text{ for all } x_i \neq x_j \in \boldsymbol{x}\}$$

for configurations $\boldsymbol{x}$ with $n(\boldsymbol{x}) > 2$.

To verify the conditions in Example 3.3, suppose that $d_G(x_i, x_j) > R$ for all distinct $x_i, x_j \in \boldsymbol{x}$ and, upon adding $u \in L \setminus \boldsymbol{x}$, that $d_G(x_i, u) > R$ for those $x_i \in \boldsymbol{x}$ that are neighbours of u in the configuration $\boldsymbol{x} \cup \{u\}$. The point x_i with minimal distance $d_G(u, x_i)$ belongs to $N(\{u\} \mid \boldsymbol{x} \cup \{u\})$, and, therefore, $d_G(x_j, u) \geq d_G(x_i, u) > R$ for all $x_j \in \boldsymbol{x}$. Thus, γ gives rise to an interaction function and a corresponding Markov density. The latter is integrable for all $\gamma > 0$ since the hard-core constraint imposes an upper bound on the number of points.

The related function $g(r) = \mathbf{1}\{r \leq R\}$ for some $R > 0$ does not satisfy the requirements for an interaction function.

3.4. The local Delaunay relation

As shown in Example 3.1, the Delaunay relation does not satisfy the consistency relations (C1) and (C2) if the graph (V, E) is not a tree. On the other hand, networks occurring in practice are seldom a tree. Therefore, we propose to employ the neighbourhood relation implicit in the graph to define a local Delaunay relation. Such a procedure is similar to that employed in image analysis for edge detection and texture analysis [9].

Definition 3.4. Let (V, E, Φ) be a graph with Euclidean edges and let L be its associated network. Define a symmetric and reflexive relation '$\sim_E$' on L as follows:

$$(i, x) \sim_E (j, y) \iff \begin{cases} \phi_i^{-1}(\partial J_i) \cap \phi_j^{-1}(\partial J_j) \neq \varnothing, & i, j \neq 0, \\ \phi_i^{-1}(\partial J_i) \cap \{y\} \neq \varnothing, & i \neq 0, j = 0, \\ \{x, y\} \in E \text{ or } x = y, & i = j = 0. \end{cases}$$

Write, for $i, j \neq 0$, '$\sim_{\boldsymbol{x}}^{i,j}$' for the Delaunay relation restricted to

$$L \cap (((\{i, 0\} \times \phi_i^{-1}(\bar{J}_i)) \cup (\{j, 0\} \times \phi_j^{-1}(\bar{J}_j))),$$

the restriction of L to at most two edges and their incident vertices, and define a symmetric reflexive relation $(i, x) \sim_z^E (j, y)$ for distinct points by

$$(i, x) \sim_z^E (j, y) \iff \begin{cases} (i, x) \sim^{i,j}_{z \cap (\{i,0\} \times \phi_i^{-1}(\bar{J}_i)) \cup (\{j,0\} \times \phi_j^{-1}(\bar{J}_j))} (j, y); (i, x) \sim_E (j, y), & i, j \neq 0, \\ (i, x) \sim^{i,i}_{z \cap (\{i,0\} \times \phi_i^{-1}(\bar{J}_i))} (0, y); (i, x) \sim_E (0, y), & i \neq 0, j = 0, \\ (0, x) \sim^{k,k}_{z \cap (\{k,0\} \times \phi_k^{-1}(\bar{J}_k))} (0, y); \{x, y\} = e_k \in E, & i = j = 0. \end{cases}$$

In words, a vertex is a $\sim_E$-neighbour of the edges it is incident with and edges are neighbours if they share a common vertex. Thus, the $\sim_E$-relation does not depend on the configuration. It does, however, crucially depend on the geometry of the graph—splitting an edge will result in a different relation. After combination with the Delaunay relation, the resulting relation '$\sim_z^E$' is configuration dependent and depends on the geometry of the graph.

To illustrate the definition, again consider Figure 1. The number of local Delaunay neighbours is one for the two topmost points, and two for the other points. As for the clique sizes, the empty set and singletons are cliques, and pairs of neighbours form a clique. There is no clique of cardinality 3.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

The main result of this section is the following.

Theorem 3.3. *Let $G = (V, E, \Phi)$ be a graph with Euclidean edges and let L be its associated network. If G does not contain any triangles then the local Delaunay relation '$\sim_{\boldsymbol{x}}^{E}$' satisfies (C1) and (C2) on the family of finite simple configurations.*

Proof. See Section 5. □

Theorem 3.3 implies the following Hammersley–Clifford theorem.

Corollary 3.1. *Suppose that the conditions of Theorem 3.3 hold, and let p be a measurable function from the set of finite simple configurations into $[0, \infty)$. Then p is a Markov function with respect to the local Delaunay relation if and only if*

$$p(\boldsymbol{x}) \propto \prod_{\boldsymbol{y} \subseteq \boldsymbol{x}} \gamma(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$$

for some interaction function $\gamma(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$ under the convention that $0^0 = 0$.

Its proof is a direct application of [3, Theorem 4.13]. Note that, other than in Theorem 3.2, cliques may contain more than two points.

Example 3.6. Let (V, E, Φ) be a graph with Euclidean edges such that (V, E) does not contain triangles and let L be the corresponding network. Consider the Strauss process defined by

$$p(\boldsymbol{x}) = \alpha \beta^{n(\boldsymbol{x})} \prod_{\{i<j\,:\, x_i, x_j \in \boldsymbol{x}\}} \gamma^{\chi(\{x_i, x_j\} \mid \boldsymbol{x})},$$

where $\gamma(\varnothing) = \alpha > 0$, $\gamma(\{x\}) = \beta > 0$, and $\gamma(\{x_1, x_2\}) = \gamma \in [0, \infty)$.

For $\gamma > 0$, $p(\boldsymbol{x}) = \alpha \beta^{n(\boldsymbol{x})} \gamma^{s^E(\boldsymbol{x})}$, where $s^E(\boldsymbol{x})$ denotes the number of point pairs $x_i \neq x_j \in \boldsymbol{x}$ that are $\sim_{\boldsymbol{x}}^{E}$-neighbours. Writing $d_{\max}$ for the maximal degree in the graph, the number of neighbours per point with respect to the local Delaunay relation is bounded by $2d_{\max}$, so $p(\boldsymbol{x})$ is integrable. It is also Markov by Corollary 3.1.

For $\gamma = 0$,

$$p(\boldsymbol{x}) = \begin{cases} \alpha \beta^{n(\boldsymbol{x})} & \text{if } n(\boldsymbol{x}) \leq 1, \\ 0 & \text{otherwise} \end{cases}$$

is subtly different from the, perhaps more natural, definition

$$\tilde{p}(\boldsymbol{x}) = \begin{cases} \alpha \beta^{n(\boldsymbol{x})} & \text{if } s^E(\boldsymbol{x}) = 0, \\ 0 & \text{otherwise.} \end{cases}$$

In words, under $\tilde{p}$, no two points are allowed to lie on the same edge or on neighbouring edges. Both functions may be normalised, but the first one is not in general Markov with respect to the local Delaunay relation. Indeed, p is defined in terms of a function $\gamma(\boldsymbol{x})$ that takes value 0 when the cardinality $n(\boldsymbol{x})$ of $\boldsymbol{x}$ exceeds 1 in accordance with property (i) of interaction functions. But then property (ii) fails, for example, for $\boldsymbol{x} = \{(i, x)\}$ and $u = (j, y)$ such that $(i, x) \not\sim_E (j, y)$.

Example 3.7. Let (V, E, Φ) be a graph with Euclidean edges such that (V, E) does not contain triangles and let L be the corresponding network. Consider the hard-core function introduced in Example 3.5 with $R > 0$ and $\gamma > 0$, and define a function γ as in that example. Then condition (ii) for an interaction function may fail with respect to the local Delaunay relation, for example, when a new point (i, x) is placed on an edge e_i that has no vertex in common with edge e_j of an existing point (j, y), but for which $d_G((i, x), (j, y)) \leq R$.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

4. Proofs of the lemmas

Proof of Lemma 3.1. We claim that $y_1 \sim_{\boldsymbol{x}} y_2$ implies that $y_1 \sim_{\boldsymbol{z}} y_2$ for any $\boldsymbol{z} \subseteq \boldsymbol{x}$ and $y_1, y_2 \in \boldsymbol{z}$. To see this, note that

$$y_1 \sim_{\boldsymbol{x}} y_2 \quad \Longleftrightarrow \quad \text{there exists } \xi \in L : \xi \in C(y_1 \mid \boldsymbol{x}) \cap C(y_2 \mid \boldsymbol{x}).$$

Since $C(y_j \mid \boldsymbol{x}) \subseteq C(y_j \mid \boldsymbol{z})$ when $\boldsymbol{z} \subseteq \boldsymbol{x}$, also $\xi \in C(y_1 \mid \boldsymbol{z}) \cap C(y_2 \mid \boldsymbol{z})$ and $y_1 \sim_{\boldsymbol{z}} y_2$.

By convention, the clique indicator function takes the value 1 for singletons and the empty set regardless of the configuration. Hence, we may focus on configurations $\boldsymbol{y}$ of cardinality at least 2.

For the first statement, suppose that $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 1$. Pick any $y_1, y_2 \in \boldsymbol{y}$. Then $y_1 \sim_{\boldsymbol{x}} y_2$ and, by the above claim, $y_1 \sim_{\boldsymbol{z}} y_2$. Since y_1 and y_2 are chosen arbitrarily, $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$.

For the second statement, if $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 0$, there exist $y_1, y_2 \in \boldsymbol{z}$ such that $y_1 \not\sim_{\boldsymbol{z}} y_2$. By the claim, also $y_1 \not\sim_{\boldsymbol{x}} y_2$; hence, $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 0$. □

Proof of Lemma 3.2. Suppose that (V, E, Φ) is a tree. Since (V, E) is a tree, there is a unique path between each pair of vertices. Hence, we may restrict ourselves to a pair of points (i, x) and (j, y) of which at least one belongs to the set $L \setminus (\{0\} \times V)$.

If $i = j$, one path runs along the edge. Since a walk does not reverse its tracks by definition, any other walk from (i, x) to (i, y) would visit at least one of the two end vertices of e_i before returning. If the same vertex were used for the return journey, the walk would not be a path. Otherwise, a cycle would be created between the end vertices in contradiction with the assumption.

If $i = 0$ and $j \neq 0$, first consider the case that $x \in \{v_j^1, v_j^2\}$. Then one path from (i, x) to (j, y) runs along the edge e_j. Any other path would have to run via the other end vertex of e_j since it cannot reverse its tracks along the edge, thus creating a cycle between the end vertices in contradiction with the assumption.

When (i, x) and (j, y) do not lie on the same closed edge, note that there is a unique path from x to v_j^1 in the graph (V, E) and, hence, a corresponding one in the labelled space L. If this path includes edge e_j, by deleting the j-labelled curve segment from y to v_j^1, we obtain a path from (i, x) to (j, y); otherwise, such a path is found by extending the path to v_j^1 with this segment. Any other path (i, x) to (j, y) would have to pass at least one of the vertices in e_j. In case this vertex would be v_j^1, by the uniqueness of paths in the tree (V, E), the path would coincide with the original construction up to the vertex and, hence, entirely. In case the vertex would be v_j^2, a cycle would be created, which cannot happen since (V, E) is a tree.

Finally, if $i \neq j$ and $i, j \neq 0$, as seen in the previous case, there is a unique path from $(0, v_i^1)$ to (j, y). If this path includes edge e_i, this yields the path from (i, x) to (j, y). Otherwise, precede by the segment along e_i from (i, x) to $(0, v_i^1)$. By the same arguments as used in the previous case, the path is unique.

Conversely, let (V, E, Φ) be such that there is a unique path between any pair of distinct points (i, x) and (j, y) in L. In particular, there is a unique path between any pair of 0-labelled vertices, and, therefore, (V, E) is a tree. □

Proof of Lemma 3.3. Since L is a tree, by Lemma 3.2, there are unique paths from y_1 to y_2 and from y_2 to y_3, say with consecutively labelled vertices

$$y_1, (0, v_1), \ldots, (0, v_p), y_2, (0, v_{p+1}), \ldots, (0, v_{p+q}), y_3,$$

$p, q \geq 0$ and no vertex replication in $v_1, \ldots, v_p$ or in $v_{p+1}, \ldots, v_{p+q}$.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

Suppose that $v_k = v_l$ for some $1 \le k \le p$ and $p+1 \le l \le p+q$. Then the paths $((0, v_k), \ldots, (0, v_p), y_2)$ and $((0, v_l), \ldots, (0, v_{p+1}), y_2)$ both connect $(0, v_k) = (0, v_l)$ and y_2 in L, and must therefore coincide. Extending, if possible, we may assume that $v_{k-1} \neq v_{l+1}$. If some earlier vertex v_i, $i \le k-1$, would be identical to v_j for $l+1 \le j \le p+q$, a cycle would be created from v_i via $v_k = v_l$ to $v_j = v_i$. Hence, the sequences $((0, v_1), \ldots, (0, v_{k-1}))$ and $((0, v_{l+1}), \ldots, (0, v_{p+q}))$ do not intersect, and the paths $(y_1, (0, v_1), \ldots, (0, v_{k-1}))$, $(y_2, (0, v_p), \ldots, (0, v_{k+1}))$ and $(y_3, (0, v_{p+q}), \ldots, (0, v_{l+1}))$ are connected at $(0, v_k) = (0, v_l)$ and, therefore, form three sides of a star provided the lengths are positive. It therefore remains to consider the cases $k = 1$, $k = p$, and $l = p+q$.

If $l = p+q$, y_3 may lie on the path from y_1 to y_2; if it does not, the curve segment from $(0, v_l)$ to y_3 forms a side of positive length. Similarly, if $k = p$, y_2 may lie on the path from y_1 to y_3; if it does not, the curve segment from $(0, v_k)$ to y_2 forms a side of positive length. Also, if $k = 1$, y_1 may lie along the path between y_2 and y_3; if it does not, the curve segment from y_1 to $(0, v_k)$ forms a side of positive length.

Finally, if there is no $1 \le k \le p$ such that $v_k = v_l$ for some $p+1 \le l \le p+q$, the unique path from y_1 to y_3 runs via y_2. □

Proof of Lemma 3.4. Suppose that $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 1$ for some $\boldsymbol{y} = \{y_1, y_2, y_3\} \subseteq \boldsymbol{x} \subset L$. By Lemma 3.1, $\chi(\boldsymbol{y} \mid \boldsymbol{y}) = 1$; hence, we need consider only $\boldsymbol{x} = \boldsymbol{y}$.

By Lemma 3.3, the elements of $\boldsymbol{y}$ either form a star or a path. First consider the case that $\boldsymbol{y}$ is a three-pointed star emanating from a centre $(0, v)$, $v \in V$. Ordering the sides according to their length, without loss of generality, suppose that $a \le b \le c$, where $a = d_G(y_1, (0, v))$, $b = d_G(y_2, (0, v))$, and $c = d_G(y_3, (0, v))$. Since, by assumption, $\boldsymbol{y}$ is in general position, at least one of the inequalities must be strict. If $a < b \le c$, since L is a tree, $C(y_2 \mid \boldsymbol{y}) \cap C(y_3 \mid \boldsymbol{y}) = \varnothing$ and $y_2 \not\sim_{\boldsymbol{y}} y_3$. The remaining case that $a = b < c$ cannot happen as it would violate the assumption that the points are in general position.

It remains to consider the case that all three elements of $\boldsymbol{y}$ lie along a path, without loss of generality from y_1 via y_2 to y_3. Then $C(y_1 \mid \boldsymbol{y}) \cap C(y_3 \mid \boldsymbol{y}) = \varnothing$, again using the tree property to ensure that there are no paths to connect y_1 and y_3 other than via y_2. Therefore, $y_1 \not\sim_{\boldsymbol{y}} y_3$. In conclusion, there cannot be a clique of cardinality 3 or larger.

To prove the second assertion, let $\boldsymbol{y} = \{y_1, y_2\}$ and write ξ for the midpoint with respect to d_G along the path between y_1 and y_2. Since L is assumed to be a tree, by Lemma 3.2, the path is unique. Moreover, as all $\phi_i \in \Phi$ are homeomorphisms, the midpoint exists and is unique. Let $\boldsymbol{x} \supseteq \boldsymbol{y}$ be some configuration in general position. If $\xi \in C(y_1 \mid \boldsymbol{x}) \cap C(y_2 \mid \boldsymbol{x})$, clearly $\chi(\boldsymbol{y} \mid \boldsymbol{x}) = 1$.

Conversely, suppose that $\boldsymbol{y}$ is a clique in $\boldsymbol{x}$. Then there exists some $\eta \in L$ such that $\eta \in C(y_1 \mid \boldsymbol{x}) \cap C(y_2 \mid \boldsymbol{x})$. By Lemma 3.3, the triple $\{\eta, y_1, y_2\}$ either forms a path or a star. In the first case, the property of equidistance to y_1 and y_2 implies that $\eta = \xi$ and the proof is complete.

Next consider the case that η and $\boldsymbol{y}$ form a three-pointed star whose sides have strictly positive length and emanate from some centre $(0, v)$. Since $d_G(\eta, y_j) = d_G(\eta, (0, v)) + d_G((0, v), y_j)$, $j = 1, 2$, by the uniqueness of the paths from η to y_j (cf. Lemma 3.2), $d_G((0, v), y_1) = d_G((0, v), y_2)$. Hence, $\xi = (0, v)$, now using the uniqueness of the path between y_1 and y_2.

It remains to show that no other point of $\boldsymbol{x}$ lies closer to ξ than the y_j, $j = 1, 2$. To this end, note that any such $x \in \boldsymbol{x}$ cannot be connected to the star by means of paths attached to different sides, as that would create a cycle.

Any $x \in \boldsymbol{x}$ connected to the endpoints y_1 or y_2 lies further from ξ than y_1 and y_2.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

For any $x \in \boldsymbol{x}$ connected to the side of η at a point other than ξ, since η is assumed to lie in $C(y_1 \mid \boldsymbol{x}) \cap C(y_2 \mid \boldsymbol{x})$,

$$d_G(\eta, x) \geq d_G(\eta, y_j) = d_G(\eta, \xi) + d_G(\xi, y_j), \qquad j = 1, 2,$$

again using the uniqueness of paths for the last equality. If the connection is at the endpoint η then $d_G(\xi, x) = d_G(\xi, \eta) + d_G(\eta, x)$; hence, by the above equation, $d_G(\xi, x) \geq 2d_G(\xi, \eta) + d_G(\xi, y_j) \geq d_G(\xi, y_j)$. If the connection is at a vertex $(0, w)$, $w \in V$,

$$\begin{aligned} d_G(\xi, x) &= d_G(\xi, (0, w)) + d_G(x, (0, w)) \\ &= d_G(\xi, (0, w)) + d_G(\eta, x) - d_G(\eta, (0, w)) \\ &\geq d_G(\xi, (0, w)) + d_G(\eta, y_j) - d_G(\eta, (0, w)) \\ &= d_G(\xi, (0, w)) + d_G(y_j, (0, w)) \\ &\geq d_G(\xi, y_j). \end{aligned}$$

Finally, any $x \in \boldsymbol{x}$ connected to the side of a y_j at ξ or some other vertex $(0, w)$ satisfies, since η is assumed to lie in $C(y_1 \mid \boldsymbol{x}) \cap C(y_2 \mid \boldsymbol{x})$,

$$d_G(\eta, x) \geq d_G(\eta, y_j) = d_G(\eta, \xi) + d_G(\xi, y_j),$$

so that $d_G(\xi, x) = d_G(\eta, x) - d_G(\eta, \xi) \geq d_G(\xi, y_j)$, and the proof is complete. □

5. Proofs of the main theorems

Proof of Theorem 3.1. First note that if $n(\boldsymbol{y}) \in \{0, 1\}$, (C1) and (C2) are automatically satisfied. The assumption that L is a tree and the restriction to configurations in general position imply, by Lemma 3.4, that $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 0$ for all $\boldsymbol{z}$ when the cardinality of $\boldsymbol{y}$ is 3 or more. Hence, (C1) and (C2) hold for $\boldsymbol{y}$ with $n(\boldsymbol{y}) \geq 3$ as well, and it suffices to consider pairs $\boldsymbol{y} = \{y_1, y_2\}$.

Condition (C1). Take $y_1, y_2 \in \boldsymbol{z} \subset L$ and $u \in L$ with $u \notin \boldsymbol{z}$, and suppose that $\chi(\{y_1, y_2\} \mid \boldsymbol{z})$ differs from $\chi(\{y_1, y_2\} \mid \boldsymbol{z} \cup \{u\})$. By Lemma 3.1, if $\chi(\{y_1, y_2\} \mid \boldsymbol{z} \cup \{u\}) = 1$ then we have $\chi(\{y_1, y_2\} \mid \boldsymbol{z}) = 1$ also, so it suffices to consider the case that $\chi(\{y_1, y_2\} \mid \boldsymbol{z} \cup \{u\}) = 0$, but $\chi(\{y_1, y_2\} \mid \boldsymbol{z}) = 1$.

Let $\tilde{y}_j$ be the point lying halfway between y_j and u along the unique path between them (cf. Lemma 3.2), and let ξ be the halfway point between y_1 and y_2. These points exist since the parametrisations are homeomorphisms. Since $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1, \xi \in C(y_1 \mid \boldsymbol{z}) \cap C(y_2 \mid \boldsymbol{z})$ by Lemma 3.4, so that, for all $z \in \boldsymbol{z}$, $d_G(\xi, z) \geq d_G(\xi, y_j)$, $j = 1, 2$. Also, by construction, $d_G(\tilde{y}_j, y_j) = d_G(\tilde{y}_j, u)$. We shall show that $\tilde{y}_j \in C(y_j \mid \boldsymbol{z} \cup \{u\}) \cap C(u \mid \boldsymbol{z} \cup \{u\})$ for $j = 1, 2$, and, therefore, $y_j \sim_{\boldsymbol{z} \cup \{u\}} u$. To this end, we must demonstrate that

$$d_G(\tilde{y}_j, z) \geq d_G(\tilde{y}_j, u) = d_G(\tilde{y}_j, y_j) \quad \text{for } z \in \boldsymbol{z} \setminus \{y_j\}. \tag{5.1}$$

By Lemma 3.3, since L is a tree, the paths between the three points u, y_1, and y_2 form either a three-pointed star whose sides emanate from $\eta = (0, v)$ with $v \in V$, or the points lie on a path.

First consider the case in which u, y_1, and y_2 lie on a path. If u were an extremity of this path, say with y_1 on the path between u and y_2, then

$$d_G(\xi, u) = d_G(\xi, y_1) + d_G(y_1, u) > d_G(\xi, y_1).$$

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

By the assumption that $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$ and Lemma 3.4, for any $z \in \boldsymbol{z}$, also $d_G(\xi, z) \geq d_G(\xi, y_1) = d_G(\xi, y_2)$, so that $\xi \in C(y_1 \mid \boldsymbol{z} \cup \{u\}) \cap C(y_2 \mid \boldsymbol{z} \cup \{u\})$, thus violating the assumption that $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u\}) = 0$. We conclude that u has to lie on the path from y_1 to y_2. Suppose that (5.1) does not hold, that is, for some j and some $z \in \boldsymbol{z} \setminus \{y_j\}$, the distance $d_G(\tilde{y}_j, z) < d_G(\tilde{y}_j, y_j)$. Then

$$d_G(\xi, z) \leq d_G(\xi, \tilde{y}_j) + d_G(\tilde{y}_j, z) < d_G(\xi, \tilde{y}_j) + d_G(\tilde{y}_j, y_j) = d_G(\xi, y_j),$$

using uniqueness of paths (Lemma 3.2). However, if $d_G(\xi, z) < d_G(\xi, y_j)$ then $\xi \notin C(y_j \mid \boldsymbol{z})$, in contradiction with Lemma 3.4.

Next suppose that the triple $\{u, y_1, y_2\}$ forms a star and, without loss of generality, that the path from y_1 to y_2 passes first ξ and then η. Since $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u\}) = 0$, the intersection of the $C(y_j \mid \boldsymbol{z} \cup \{u\})$, $j = 1, 2$, is empty and, in particular, does not contain ξ. Hence,

$$d_G(\xi, u) < d_G(\xi, y_1) = d_G(\xi, y_2).$$

Therefore,

$$d_G(\eta, u) < d_G(\eta, y_2) \leq d_G(\eta, y_1),$$

with equality only if $\eta = \xi$. Consequently, $\tilde{y}_j$ lies on the side of y_j for each $j = 1, 2$. Now, if $d_G(\tilde{y}_2, z) < d_G(\tilde{y}_2, y_2)$ for some $z \in \boldsymbol{z} \setminus \{y_2\}$, since $\tilde{y}_2$ lies on the path from ξ to y_2 via η, then

$$d_G(\xi, z) \leq d_G(\xi, \tilde{y}_2) + d_G(\tilde{y}_2, z) < d_G(\xi, \tilde{y}_2) + d_G(\tilde{y}_2, y_2) = d_G(\xi, y_2),$$

in contradiction with the assumption that $\xi \in C(y_2 \mid \boldsymbol{z})$. Similarly, for y_1, recalling that $d_G(\xi, u) < d_G(\xi, y_1)$, the point $\tilde{y}_1$ lies on the path between ξ and y_1. Hence, if $d_G(\tilde{y}_1, z) < d_G(\tilde{y}_1, y_1)$ then

$$d_G(\xi, z) \leq d_G(\xi, \tilde{y}_1) + d_G(\tilde{y}_1, z) < d_G(\xi, \tilde{y}_1) + d_G(\tilde{y}_1, y_1) = d_G(\xi, y_1),$$

in contradiction with the assumption that $\xi \in C(y_1 \mid \boldsymbol{z})$. Therefore, (5.1) cannot be violated and (C1) holds.

Condition (C2). Take $y_1, y_2 \in \boldsymbol{z} \subset L$ and $u_1, u_2 \in L$ such that $u_1, u_2 \notin \boldsymbol{z}$ with $u_1 \not\sim_{\boldsymbol{x}} u_2$, where $\boldsymbol{x} = \boldsymbol{z} \cup \{u_1, u_2\}$. Write $\boldsymbol{y} = \{y_1, y_2\}$. Lemma 3.1 implies that if $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 0$, the same is true upon adding points to $\boldsymbol{z}$ and (C2) holds in this case. Therefore, it suffices to consider the case that $\chi(\{y_1, y_2\} \mid \boldsymbol{z}) = 1$. The same lemma implies that if $\chi(\{y_1, y_2\} \mid \boldsymbol{z} \cup \{u_1, u_2\}) = 1$, this remains true when deleting points from $\boldsymbol{x}$ and (C2) holds. The only case left to consider is that when $\chi(\{y_1, y_2\} \mid \boldsymbol{z}) = 1$ and $\chi(\{y_1, y_2\} \mid \boldsymbol{z} \cup \{u_1, u_2\}) = 0$. We must show that exactly one of $\chi(\{y_1, y_2\} \mid \boldsymbol{z} \cup \{u_1\})$ and $\chi(\{y_1, y_2\} \mid \boldsymbol{z} \cup \{u_2\})$ takes the value 1 and will do so by contradiction.

Let $\eta \in L$ be the point that lies halfway between u_1 and u_2 along the unique path between them (cf. Lemma 3.2), and write ξ for the halfway point between y_1 and y_2. These points exist since the parametrisations are homeomorphisms.

Suppose that $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1\}) = 0 = \chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_2\})$. By the proof of condition (C1) above, the triple $\{u_1, y_1, y_2\}$ forms a direct path with u_1 on the path from y_1 to y_2, or a three-pointed star. In either case,

$$d_G(\xi, u_1) < d_G(\xi, y_1) = d_G(\xi, y_2).$$

The same is true for the triple $\{u_2, y_1, y_2\}$. Therefore, the ensemble can be seen as a path from y_1 to y_2 that passes points leading off to u_1 and to u_2 if we include the degenerate cases of a branch consisting of the single point u_1 or u_2, respectively. Without loss of generality, suppose that the order is y_1, then the side(s) of u_1 and u_2, and finally y_2. We claim that such an arrangement would imply that $u_1 \sim_{\boldsymbol{x}} u_2$, in contradiction with the assumption.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

To prove the claim, we show that $\eta \in C(u_1 \mid \boldsymbol{z} \cup \{u_1, u_2\}) \cap C(u_2 \mid \boldsymbol{z} \cup \{u_1, u_2\})$, that is, $d_G(\eta, u_1) = d_G(\eta, u_2) \leq d_G(\eta, z)$ for all $z \in \boldsymbol{z}$. Suppose otherwise. Then, for some $z \in \boldsymbol{z}$, possibly y_1 or y_2, $d_G(\eta, z) < d_G(\eta, u_1) = d_G(\eta, u_2)$ and, therefore,

$$d_G(\xi, z) \leq d_G(\xi, \eta) + d_G(\eta, z) < d_G(\xi, \eta) + d_G(\eta, u_1) = d_G(\xi, \eta) + d_G(\eta, u_2).$$

The right-hand side is equal to either $d_G(\xi, u_1)$ or $d_G(\xi, u_2)$ and, upon recalling that both are strictly smaller than $d_G(\xi, y_j)$, we obtain $d_G(\xi, z) < d_G(\xi, y_1) = d_G(\xi, y_2)$, in contradiction with the assumption that $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$, that is, $\xi \in C(y_j \mid \boldsymbol{z})$ for $j = 1, 2$ (Lemma 3.4). Hence, $C(u_1 \mid \boldsymbol{z} \cup \{u_1, u_2\}) \cap C(u_2 \mid \boldsymbol{z} \cup \{u_1, u_2\})$ contains η, implying that $u_1 \sim_{\boldsymbol{x}} u_2$ in contradiction with the assumption.

Finally, suppose that $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1\}) = 1 = \chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_2\})$. By Lemma 3.4, we have $\xi \in C(y_1 \mid \boldsymbol{z} \cup \{u_1\}) \cap C(y_2 \mid \boldsymbol{z} \cup \{u_1\})$ and, therefore, $d_G(\xi, u_1) \geq d_G(\xi, y_1) = d_G(\xi, y_2)$. Similarly, $d_G(\xi, u_2) \geq d_G(\xi, y_1) = d_G(\xi, y_2)$. By assumption, $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1, u_2\}) = 0$, which means that $C(y_1 \mid \boldsymbol{z} \cup \{u_1, u_2\}) \cap C(y_2 \mid \boldsymbol{z} \cup \{u_1, u_2\}) = \varnothing$ and, in particular, does not contain ξ. Therefore, recalling the assumption that $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$, we have $\min(d_G(\xi, u_1), d_G(\xi, u_2)) < d_G(\xi, y_1) = d_G(\xi, y_2)$, a contradiction.

In conclusion, exactly one of $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1\})$ and $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_2\})$ takes the value 1. □

Proof of Theorem 3.2. Suppose that p is a Markov function. By Theorem 3.1, the Delaunay relation satisfies (C1) and (C2), and [3, Theorem 4.13] implies that p can be factorised as

$$p(\boldsymbol{x}) = \prod_{\boldsymbol{y} \subseteq \boldsymbol{x}} \gamma(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$$

for some interaction function $\gamma(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$ under the convention that $0^0 = 0$. Since the cliques have cardinality at most 2 (cf. Lemma 3.4), the factorisation reduces to

$$p(\boldsymbol{x}) = \gamma(\varnothing) \prod_i \gamma(\{x_i\}) \prod_{i<j} \gamma(\{x_i, x_j\})^{\chi(\{x_i, x_j\} \mid \boldsymbol{x})} \Bigg(\prod_{\{\boldsymbol{y} \subseteq \boldsymbol{x} \colon n(\boldsymbol{y}) > 2\}} \gamma(\boldsymbol{y}) \Bigg)^0.$$

Note that $p(\boldsymbol{x}) > 0$ if and only if $\gamma(\boldsymbol{y}) > 0$ for all $\boldsymbol{y} \subseteq \boldsymbol{x}$. We claim that it suffices to consider only subsets of cardinality at most 2, that is, $p(\boldsymbol{x}) = 0$ if and only if there exists some $\boldsymbol{y} \subseteq \boldsymbol{x}$ with $n(\boldsymbol{y}) \leq 2$ such that $\gamma(\boldsymbol{y}) = 0$. Then

$$p(\boldsymbol{x}) = \gamma(\varnothing) \prod_i \gamma(\{x_i\}) \prod_{\{i<j \colon x_i \sim_{\boldsymbol{x}} x_j\}} \gamma(\{x_i, x_j\}),$$

unless $\gamma(\{x_i, x_j\}) = 0$ for some $x_i, x_j \in \boldsymbol{x}$ in which case $p(\boldsymbol{x}) = 0$. Moreover,

$$p(\boldsymbol{x}) = \prod_{\boldsymbol{y} \subseteq \boldsymbol{x}} \tilde{\gamma}(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$$

for $\tilde{\gamma}(\boldsymbol{y})$ equal to $\gamma(\boldsymbol{y})$ when $n(\boldsymbol{y}) \leq 2$ and $\tilde{\gamma}(\boldsymbol{y}) = \mathbf{1}\{\gamma(\boldsymbol{z}) > 0 \text{ for all } \boldsymbol{z} \subset \boldsymbol{y} \text{ with } n(\boldsymbol{z}) \leq 2\}$ otherwise. The function $\tilde{\gamma}(\boldsymbol{y})^{\chi(\boldsymbol{y} \mid \boldsymbol{x})}$ is an interaction function. Indeed, since, as an interaction function, $\gamma(\boldsymbol{x}) > 0$ if and only if $\gamma(\boldsymbol{y}) > 0$ for all $\boldsymbol{y} \subseteq \boldsymbol{x}$, so, according to the claim, if and only if $\tilde{\gamma}(\boldsymbol{x}) > 0$, properties (i) and (ii) carry over from γ to $\tilde{\gamma}$.

To prove the claim, we will inductively show that, when $\gamma(\boldsymbol{y}) > 0$ for all subsets with $n(\boldsymbol{y}) \leq k$, $k \geq 2$, the same is true for sets of cardinality $k + 1$. Suppose otherwise. Then $\gamma(\boldsymbol{y} \cup \{u\}) = 0$ for some $\boldsymbol{y}$ with $n(\boldsymbol{y}) = k$ and some $u \notin \boldsymbol{y}$. As γ is an interaction function, if $\gamma(N(\{u\} \mid \boldsymbol{y} \cup \{u\}))$ is strictly positive, also $\gamma(\boldsymbol{y} \cup \{u\}) > 0$, in contradiction with the assumption. Hence, $\gamma(N(\{u\} \mid \boldsymbol{y} \cup \{u\})) = 0$. Therefore, the cardinality of $N(\{u\} \mid \boldsymbol{y} \cup \{u\})$

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

must be $k+1$. Hence, $N(\{u\} \mid \boldsymbol{y} \cup \{u\}) = \boldsymbol{y} \cup \{u\}$, that is, $y_i \sim_{\boldsymbol{y}\cup\{u\}} u$ for all $y_i \in \boldsymbol{y}$. Because $k+1 \geq 3$, by Lemma 3.4, the set $\boldsymbol{y} \cup \{u\}$ is no clique and there exist $y_1, y_2 \in \boldsymbol{y}$ for which $y_1 \not\sim_{\boldsymbol{y}\cup\{u\}} y_2$. Therefore, $y_2 \notin N(\{y_1\} \mid \boldsymbol{y} \cup \{u\})$ and, hence, $\gamma(N(\{y_1\} \mid \boldsymbol{y} \cup \{u\})) > 0$ as well as $\gamma(\{u\} \cup \boldsymbol{y} \setminus \{y_1\}) > 0$ in view of their cardinalities. Property (ii) implies that $\gamma(\boldsymbol{y} \cup \{u\}) > 0$ and we arrive at a contradiction.

Conversely, any function of the specified form is a Markov function. To see this, suppose that $p(\boldsymbol{x}) > 0$. Then $\gamma(\boldsymbol{y}) > 0$ for all $\boldsymbol{y} \subseteq \boldsymbol{x}$. Hence, $p(\boldsymbol{y}) > 0$ for all $\boldsymbol{y} \subseteq \boldsymbol{x}$. Furthermore, the ratio $p(\boldsymbol{x} \cup \{\xi\})/p(\boldsymbol{x})$ can be written as

$$\gamma(\{\xi\})\, \mathbf{1}\{\gamma(\boldsymbol{x} \cup \{\xi\}) > 0\} \prod_{\{i\,:\, x_i \sim_{\boldsymbol{x}\cup\{\xi\}} \xi\}} \gamma(\{x_i, \xi\}) \frac{\prod_{\{i<j\,:\, x_i \sim_{\boldsymbol{x}\cup\{\xi\}} x_j\}} \gamma(\{x_i, x_j\})}{\prod_{\{i<j\,:\, x_i \sim_{\boldsymbol{x}} x_j\}} \gamma(\{x_i, x_j\})}.$$

By Theorem 3.1, the Delaunay relation satisfies condition (C1), so that, if $\chi(\{x_i, x_j\} \mid \boldsymbol{x}) \neq \chi(\{x_i, x_j\} \mid \boldsymbol{x} \cup \{\xi\})$, both x_i and x_j belong to the neighbourhood $N(\{\xi\} \mid \boldsymbol{x} \cup \{\xi\})$. Hence, the last term in the ratio of interest depends only on points in the configuration $\boldsymbol{x} \cap N(\{\xi\} \mid \boldsymbol{x} \cup \{\xi\})$ and on the relations '$\sim_{\boldsymbol{x}}$', '$\sim_{\boldsymbol{x}\cup\{\xi\}}$' restricted to this configuration. The product of $\gamma(\{x_i, \xi\})$ over points in the configuration $\boldsymbol{x} \cap N(\{\xi\} \mid \boldsymbol{x} \cup \{\xi\})$ depends on this configuration and on ξ, $\gamma(\{\xi\})$ is a function of ξ, and, finally, by property (ii) of γ, to verify that $\gamma(\boldsymbol{x} \cup \{\xi\}) > 0$, it suffices to verify that $\gamma(N(\{\xi\} \mid \boldsymbol{x} \cup \{\xi\})) > 0$, which depends on ξ and $\boldsymbol{x} \cap N(\{\xi\} \mid \boldsymbol{x} \cup \{\xi\})$. □

Proof of Theorem 3.3. Let $\boldsymbol{y} \subseteq \boldsymbol{z} \subset L$ and $u, v \in L$ be such that $u, v \notin \boldsymbol{z}$ and $u \not\sim_{\boldsymbol{z}\cup\{u,v\}} v$.

We first observe that cliques in '$\sim_E$' consist of points lying either on a single closed edge, on a triangle of edges, or on two or more closed edges emanating from a single vertex. The no-triangle assumption excludes the second case.

The Delaunay relation restricted to a pair of such edges coincides with the sequential neighbourhood relation on the edges. In other words, consecutive points on a single edge are each other's nearest neighbours; also the point on one of the edges that is closest to the vertex that joins the two edges, if it exists, is a nearest neighbour of the point closest to that vertex on the other edge, and no other points are nearest neighbours. Cliques in the combined relation '$\sim_{\boldsymbol{z}}^{E}$' are therefore either empty, consist of a single point, of two consecutive points on a single edge, or of points on different edges that are closest on their edge to the central vertex from which all edges emanate. The clique size is therefore at most the degree of the central vertex.

Condition (C1). If $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 0$, there exists a pair $y_1 \neq y_2 \in \boldsymbol{y}$ for which $y_1 \not\sim_{\boldsymbol{z}}^{E} y_2$. Then either y_1 and y_2 lie on edges that are not adjacent, in which case $y_1 \not\sim_{\boldsymbol{z}\cup\{u\}}^{E} y_2$ and $\chi(\boldsymbol{y} \mid \boldsymbol{z}\cup\{u\}) = 0$, or y_1 and y_2 lie on related edges, but y_1 and y_2 are not sequential neighbours in $\boldsymbol{z}$ restricted to their edge(s). The addition of u cannot make them sequential neighbours, so $\chi(\boldsymbol{y} \mid \boldsymbol{z}\cup\{u\}) = 0$.

Suppose therefore that $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$, but that there exists a pair $y_1 \neq y_2 \in \boldsymbol{y}$ for which $y_1 \not\sim_{\boldsymbol{z}\cup\{u\}}^{E} y_2$. Then y_1 and y_2 must lie on $\sim_E$-related edges (either a single edge, or two adjacent edges) and be consecutive in $\boldsymbol{z}$ but not in $\boldsymbol{z} \cup \{u\}$ on these edge(s). This can happen only if u lies in between y_1 and y_2, making y_1 and y_2 both $\sim_{\boldsymbol{z}\cup\{u\}}^{E}$-neighbours of u. Should $\boldsymbol{y}$ contain additional points, say y_3, by the general remarks, y_3 lies on a different edge closest to the central vertex, as do y_1 and y_2 on their respective edges. Therefore, y_3 is a neighbour of u with respect to '$\sim_{\boldsymbol{z}\cup\{u\}}^{E}$'.

Condition (C2). As shown when proving (C1), if $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 0$, $\boldsymbol{y}$ cannot be a clique in configurations with more points. If $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1, u_2\}) = 1$, all pairs of points in $\boldsymbol{y}$ lie on adjacent edges and are sequential neighbours in the set $\boldsymbol{z} \cup \{u_1, u_2\}$) restricted to their edges. The same remains true when u_1 and u_2 are removed, so that $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$. Hence, it remains

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press

to consider the case when $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$ but $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1, u_2\}) = 0$. In this case, the points of $\boldsymbol{y}$ must lie on a single edge or on a number of edges that emanate from a single vertex $(0, w) \in L,\ w \in V$.

Now, if u_1 and u_2 do not lie on any of these $\boldsymbol{y}$-edges, $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1, u_2\}) = 1$, in contradiction with the assumptions.

If exactly one of u_1 and u_2 does not lie on any of the $\boldsymbol{y}$-edges, without loss of generality, assume that u_1 does not lie on any of the $\boldsymbol{y}$-edges and that u_2 does lie on some $\boldsymbol{y}$-edge. Then $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1\}) = \chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$. Moreover, since, by assumption, $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_1, u_2\}) = 0$, there is a pair $y_1, y_2 \in \boldsymbol{y}$ that are not consecutive in the configuration $\boldsymbol{z} \cup \{u_1, u_2\}$ restricted to the edge or edges on which y_1 and y_2 lie. Since they are adjacent in the configuration $\boldsymbol{z}$, u_2 must lie in between y_1 and y_2, which implies that $\chi(\boldsymbol{y} \mid \boldsymbol{z} \cup \{u_2\}) = 0$ in accordance with (C2).

Finally, suppose that both u_1 and u_2 lie on the $\boldsymbol{y}$-edges that emanate from $(0, w)$. If $\boldsymbol{y} = \{y_1, y_2\}$ consists of two consecutive points in $\boldsymbol{z}$, since y_1 and y_2 are no longer consecutive in the configuration $\boldsymbol{z} \cup \{u_1, u_2\}$, they must be separated by either u_1 or u_2, but not by both, since, by assumption, u_1 and u_2 are not sequential neighbours in $\boldsymbol{z} \cup \{u_1, u_2\}$. Hence, (C2) holds. If $\boldsymbol{y}$ consists of more than two points $\boldsymbol{y} = \{y_1, \ldots, y_k\}$, the assumption that $\chi(\boldsymbol{y} \mid \boldsymbol{z}) = 1$ implies that the $y_j,\ j = 1, \ldots, k$, must lie on different edges $e_1, \ldots, e_k$ emanating from w and no points of $\boldsymbol{z}$ lie between y_i and w. Since u_1 and u_2 are not sequential neighbours, they cannot both lie between some y_i and $(0, w)$; one of them, however, must, since the clique indicator function of $\boldsymbol{y}$ in $\boldsymbol{z} \cup \{u_1, u_2\}$ takes the value 0. Consequently, also in this case (C2) is seen to hold. □

References

[1] Anderes, E., Møller, J. and Rasmussen, J. G. (2016). Second-order pseudo-stationary random fields and point processes on graphs and their edges. *AU Workshop Stoch. Geometry, Stereology, Appl.* conf. presentation. Available at http://www.lebesgue.fr/sites/default/files/attach/GraphWithEdgesSHORT.pdf.

[2] Anderes, E., Møller, J. and Rasmussen, J. G. (2017). Isotropic covariance functions on graphs and their edges. Preprint. Available at http://arxiv.org/abs/1710.01295v1.

[3] Baddeley, A. and Møller, J. (1989). Nearest-neighbour Markov point processes and random sets. *Internat. Statist. Rev.* **57,** 89–121.

[4] Baddeley, A. J., van Lieshout, M. N. M. and Møller, J. (1996). Markov properties of cluster processes. *Adv. Appl. Prob.* **28,** 346–355.

[5] Baddeley, A., Nair, G., Rakshit, S. and McSwiggan, G. (2017). "Stationary" point processes are uncommon on linear networks. *Stat* **6,** 68–78.

[6] Bondy, A. and Murty, M. R. (2008). *Graph Theory*. Springer, London.

[7] Daley, D. J. and Vere–Jones, D. (2003). *An Introduction to the Theory of Point Processes*, Vol. I, 2nd edn. Springer, New York.

[8] Daley, D. J. and Vere–Jones, D. (2008). *An Introduction to the Theory of Point Processes*, Vol. II, 2nd edn. Springer, New York.

[9] Geman, D., Geman, S., Graffigne, C. and Dong, P. (1990). Boundary detection by constrained optimization. *IEEE Trans. Pattern Anal. Mach. Intellig.* **12,** 609–628.

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

[11] Jónsdóttir, K. Ý., Hahn, U. and Jensen, E. B. V. (2004). Inhomogeneous spatial point processes, with a view to spatio-temporal modelling. In *Spatial Point Process Modelling and Its Applications* (Lecture Notes Statist. **185**), eds A. Baddeley *et al.*, Springer, New York, pp. 131–136.

[12] Okabe, A. and Sugihara, K. (2012). *Spatial Analysis Along Networks. Statistical and Computational Methods.* John Wiley, Chichester.

[13] Rakshit, S., Nair, G. and Baddeley, A. (2017). Second-order analysis of point patterns on a network using any distance metric. *Spatial Statist.* **22,** 129–154.

[14] Ripley, B. D. and Kelly, F. P. (1977). Markov point processes. *J. London Math. Soc.* **15,** 188–192.

[15] Van Lieshout, M. N. M. (2000). *Markov Point Processes and Their Applications*. Imperial College Press.

https://doi.org/10.1017/apr.2018.60 Published online by Cambridge University Press