Elect. Comm. in Probab. **8** (2003), 17–27

ELECTRONIC COMMUNICATIONS in PROBABILITY

TREES AND MATCHINGS FROM POINT PROCESSES

ALEXANDER E. HOLROYD[1]
Department of Mathematics, UC Berkeley, CA 94720-3840, USA
email: `holroyd@math.berkeley.edu`

YUVAL PERES[2]
Departments of Statistics and Mathematics, UC Berkeley, CA 94720-3860, USA
email: `peres@stat.berkeley.edu`

Submitted 29 November 2002 , *accepted in final form* 5 February 2003

AMS 2000 Subject classification: Primary 60G55; Secondary 60K35
Keywords: Poisson process, point process, random tree, random matching, minimal spanning forest

Abstract

A **factor graph** of a point process is a graph whose vertices are the points of the process, and which is constructed from the process in a deterministic isometry-invariant way. We prove that the d-dimensional Poisson process has a one-ended tree as a factor graph. This implies that the Poisson points can be given an ordering isomorphic to the usual ordering of the integers in a deterministic isometry-invariant way. For $d \geq 4$ our result answers a question posed by Ferrari, Landim and Thorisson [7]. We prove also that any isometry-invariant ergodic point process of finite intensity in Euclidean or hyperbolic space has a perfect matching as a factor graph provided all the inter-point distances are distinct.

1 Introduction

Let M be an isometry-invariant point process on $\mathbb{R}^d$, viewed as a random Borel measure. We assume throughout that all point processes are simple and of finite intensity. The **support** of M is $[M] = \{x \in \mathbb{R}^d : M(\{x\}) = 1\}$, and (random) elements of $[M]$ are called M-points. By a **factor graph** of M we mean a random (directed or undirected) graph G whose vertex set equals $[M]$, such that G is a deterministic function of M, and such that the joint distribution of M and G is invariant under isometries of $\mathbb{R}^d$. (We give a more formal definition at the end of the introduction).

A graph is **locally finite** if no vertex has infinite degree. A graph is a **tree** if it is connected and has no cycles. The number of **ends** of a tree is the number of distinct singly infinite self-avoiding paths from any one vertex. A **directed doubly infinite path** is a directed

[1]RESEARCH PARTIALLY SUPPORTED BY NSF GRANT DMS–0072398

[2]RESEARCH FUNDED IN PART BY NSF GRANT DMS-0104073 AND A MILLER PROFESSORSHIP AT UC BERKELEY

graph isomorphic to the graph with vertex set $\mathbb{Z}$ and a directed edge from n to $n+1$ for each n.

Theorem 1 *Let M be a Poisson point process on $\mathbb{R}^d$.*

(i) M has a factor graph which is almost surely a locally finite one-ended tree.

(ii) M has a factor graph which is almost surely a directed doubly infinite path.

Let $\|\cdot\|$ denote the Euclidean norm, and let $B(r) = \{x \in \mathbb{R}^d : \|x\| < r\}$ be the ball of radius r. A point process M is said to be **non-equidistant** if there do not exist M-points w, x, y, z with $\{w, x\} \neq \{y, z\}$ and $\|w - x\| = \|y - z\| > 0$. A graph is called a **matching** if every vertex has degree 0 or 1, and a **perfect matching** if every vertex has degree 1.

Theorem 2 *Let M be a non-equidistant point process in $\mathbb{R}^d$ which is invariant and ergodic under isometries.*

(i) M has a factor graph in which almost surely each component is a locally finite one-ended tree.

(ii) M has a factor graph in which almost surely each component is a directed doubly infinite path.

(iii) M has a factor graph which is almost surely a perfect matching.

Ferrari, Landim and Thorisson [7] proved by a different method that the d-dimensional Poisson process has a one-ended tree as a *translation-invariant* factor graph for $d \leq 3$, and asked whether this holds for $d \geq 4$. Theorem 1 (i) establishes the stronger isometry-invariant statement for all d. We will deduce the (ii) part of Theorem 1 from the (i) part. This implication was also noted in [7]. Theorem 1 (ii) is clearly equivalent to the assertion that the M-points can be given an ordering isomorphic to the usual ordering of $\mathbb{Z}$ in a deterministic isometry-invariant way. Such orderings have connections with the notion of point-stationarity, while the construction of perfect matchings as in Theorem 2 (iii) has connections with Palm processes. See [7],[13] for more details.

The main novelty of Theorem 2 lies in the generality of the point process. For the Poisson process, Theorem 2 may be proved by relatively simple constructions, including one which we discuss in Section 4.

The minimal spanning forest is a natural factor graph. However, in general it is unknown how many components it has, and how many ends the components have. (Partial answers are provided in [3],[4]).

Not every isometry-invariant ergodic point process on $\mathbb{R}^d$ has a tree as a graph factor. For example, consider the point set obtained by applying a uniform random translation and a uniform random rotation to $\mathbb{Z}^d$. This process has no perfect matching as a factor graph, and no factor graph in which every component is an infinite tree if $d \geq 2$.

The condition of non-equidistance in Theorem 2 may be relaxed to the condition that the symmetry group of $[M]$ is almost surely trivial. This is explained at the end of Section 3.

Theorems 1 extends to other amenable spaces in place of $\mathbb{R}^d$, with the same proof. In non-amenable spaces, no invariant one-ended tree exists [1],[5], and in certain spaces no invariant tree exists [2],[11]. Theorem 2 extends to other spaces as follows.

Theorem 3 *Let Λ be a locally compact metric space and let Γ be a transitive unimodular group of isometries of Λ. Let ν be a Γ-invariant Borel measure on Λ which is finite on bounded sets. Let M be a non-equidistant point process of finite intensity with respect to ν, which is invariant and ergodic under Γ. Then statements (i)–(iii) of Theorem 2 hold.*

The proof of Theorem 2 extends directly to Theorem 3 in the case of amenable spaces, but breaks down for non-amenable spaces such as hyperbolic space. We therefore give a different argument which is valid in amenable and non-amenable settings. The argument is based on the simple and appealing idea of iteratively matching mutually nearest neighbors, and is of interest even in the case of $\mathbb{R}^d$. For the Poisson process, the resulting matching is the unique stable matching in the sense of Gale and Shapley [8] - see Section 4 for the definition.
Theorems 1,2,3 are proved Sections 2,3,4 respectively.

Formal definition of a factor graph. A **factor graph** of a point process M on $\mathbb{R}^d$ is a mapping, commuting with isometries, that assigns to a countable set $V \subseteq \mathbb{R}^d$ a graph $G = G(V)$ with vertex set V, defined for almost all V with respect to the law of $[M]$. Recall that the law of M is a probability distribution on the space $\mathcal{M}$ of Borel measures on $\mathbb{R}^d$ with the corresponding completed Borel σ-algebra $\mathcal{B}$. We will show that the events considered in Theorems 1–3 contain Borel events of probability one, hence they are $\mathcal{B}$-measurable. Similar remarks apply to clumpings as defined in Section 2.

2 Poisson Trees

Lemma 4 *Let M be an isometry-invariant point process on $\mathbb{R}^d$, and let G be a factor graph all of whose components are trees with at most r ends for some constant $r < \infty$. Then G is locally finite almost surely.*

In this section we will need only the special case $r = 1$. The proof of Lemma 4 uses a version of the "mass transport principle", Lemma 5. See [5],[6],[9] for background. A **mass transport** is a non-negative measurable function $T(x, y, M)$ on $\mathbb{R}^d \times \mathbb{R}^d \times \mathcal{M}$ (where $\mathcal{M}$ is the space of Borel measures on $\mathbb{R}^d$) which is non-zero only when x, y are M-points, and which is isometry-invariant in the sense that $T(\theta x, \theta y, \theta M) = T(x, y, M)$ for any isometry θ of $\mathbb{R}^d$. We think of $T(x, y, M)$ as the mass sent from x to y when the point configuration is M. For Borel sets $A, B \subseteq \mathbb{R}^d$ we write

$$t(A, B) = \mathbf{E} \sum_{\substack{x \in A \cap [M],\\ y \in B \cap [M]}} T(x, y, M)$$

for the expected total mass sent from A to B, and we write $K = [0,1)^d \subseteq \mathbb{R}^d$ for the unit cube.

Lemma 5 *Let M be an isometry-invariant point process on $\mathbb{R}^d$, and let t be a mass transport. We have*

$$t(K, \mathbb{R}^d) = t(\mathbb{R}^d, K).$$

Proof. The isometry-invariance of T implies isometry-invariance of t. Hence, using monotone convergence to exchange expectations and sums, we have

$$t(K, \mathbb{R}^d) = \sum_{z \in \mathbb{Z}^d} t(K, K + z) = \sum_{z \in \mathbb{Z}^d} t(K - z, K) = t(\mathbb{R}^d, K).$$

□

Proof of Lemma 4. Consider the mass transport in which $T(x, y, M) = 1$ whenever G has a singly infinite self-avoiding path from x which includes the edge (x, y), and $T(x, y, M) = 0$

otherwise. Thus each vertex sends one unit of mass to each of at most r of its neighbors, while for every edge (x, y) in an infinite component, either x sends one unit to y or y sends one unit to x, or both. Applying Lemma 5, the assumption of finite intensity implies that the expected total mass received by all M-points in K is finite, so in particular it follows that no vertex can have infinite degree. $\square$

A **clumping** of a point process M is a sequence $\mathcal{P}_1, \mathcal{P}_2, \ldots$ of successively coarser partitions of $[M]$, defined from M in a deterministic isometry-invariant way. (More precisely, a clumping is a mapping, commuting with isometries, that assigns to a countable set $V \subseteq \mathbb{R}^d$ a sequence of successively coarser partitions of V, defined for almost all V with respect to the law of $[M]$). We call the elements of the partitions **clumps**, and we call the clumping **locally finite** if all clumps are finite. A **component** of a clumping is subset of $[M]$ which is the limit of some increasing sequence of clumps $A_1, A_2, \ldots$, where $A_k \in \mathcal{P}_k$. A clumping is **connected** if it has only one component.

Lemma 6 *Let M be a non-equidistant isometry-invariant point process on $\mathbb{R}^d$. If M has a clumping which is almost surely connected and locally finite then it has a factor graph which is a almost surely a locally finite one-ended tree.*

Proof. Let $\mathcal{P}_1, \mathcal{P}_2, \ldots$ be a connected locally finite clumping. We will define for each clump A a distinguished element $x \in A$ called the **leader** of A, and we will construct a factor graph G. We do this inductively as follows. First consider a clump $A \in \mathcal{P}_1$. Choose the leader x of A as follows. If $|A| = 1$, let x be the unique element. Otherwise, let x', x'' be the unique pair of M-points in A whose Euclidean distance is minimum. Then let x be the one of x', x'' which minimizes $\min\{\|x - y\| : y \in [M] \setminus \{x', x''\}\}$. The fact that M is a non-equidistant point process ensures that all the minima involved are unique almost surely. Let G have an edge from x to each of the other elements of A. Apply the same construction to every $A \in \mathcal{P}_1$.
Now suppose that leaders have been defined for all clumps in the partitions $\mathcal{P}_1, \ldots, \mathcal{P}_{k-1}$. Consider a clump $A \in \mathcal{P}_k$, let $B_1, \ldots, B_m$ be the clumps of $\mathcal{P}_{k-1}$ which are subsets of A, and let $y_1, \ldots, y_k$ be their respective leaders. Choose the leader x of A from among $y_1, \ldots, y_k$ in a deterministic isometry-invariant way by the same procedure as above. Let G have an edge from x to each $y_i \neq x$. Apply the same construction to every $A \in \mathcal{P}_k$.
By the construction G is clearly a factor graph. Also, if A is any clump, then the subgraph of G induced by A is a finite tree, and only the leader of A has any edges to vertices outside A. This implies immediately that all components of G are one-ended trees or finite trees. But G is connected since the clumping is connected, hence G is a one-ended tree. Finally G is locally finite by Lemma 4. $\square$

Lemma 7 *Let M be a non-equidistant isometry-invariant point process on $\mathbb{R}^d$. If M has a factor graph which is almost surely a locally finite one-ended tree then it has a factor graph which is almost surely a directed doubly infinite path.*

Proof. Let G be such a one-ended tree. The **parent** of a vertex x is the unique neighbor of x which lies on the infinite self-avoiding path from x. If y is the parent of x then x is a **child** of y. Vertices are **siblings** if they have the same parent, and x is a **descendant** of y if the infinite self-avoiding path from x passes through y. First order the children of each vertex in order of distance to the parent. Then order all the vertices G according to depth-first search; that is, each vertex precedes all its children, while if x precedes its sibling y then all descendants of x precede all descendants of y. (See [7] for another description of this construction). $\square$

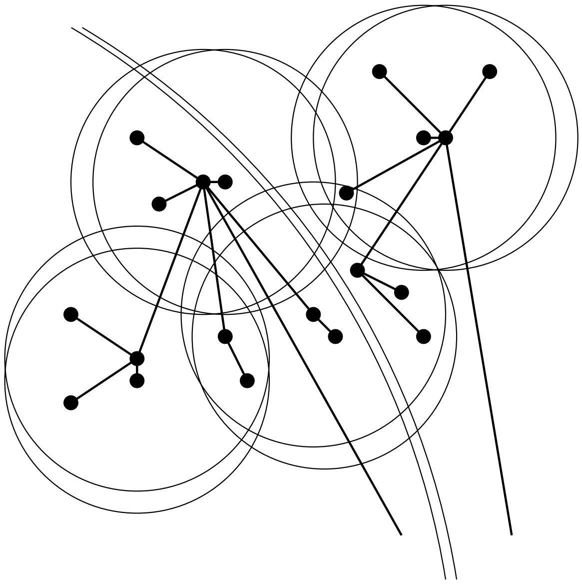

Figure 1: An illustration of the proof of Theorem 1: seeds, cutters and blobs.

PROOF OF THEOREM 1. Suppose without loss of generality that M has intensity 1. Throughout we will use $C_1, C_2, \ldots$ to denote constants in $(0, \infty)$ depending only on d. We will construct a connected clumping of M; then the theorem follows immediately from Lemmas 4,6,7.
Our argument is based on a construction in [5]. For each $k \geq 1$, let $a_k = \exp\left[-k\left(1 - \frac{1}{2d}\right)\right]$. Call an M-point x a ***k*-seed** if there is another M-point within Euclidean distance a_k of x. Clearly, the k-seeds form an isometry-invariant point process of intensity

$$\lambda_k := 1 \times \mathbf{P}\left[M(B(a_k)) \geq 1\right] = 1 - e^{-C_1 a_k^d} \leq C_1 a_k^d = C_1 e^{-k(d-1/2)}. \tag{1}$$

Now let $r_k = e^k$, and define a ***k*-cutter** to be any subset of $\mathbb{R}^d$ of the form $\{y : \|y - x\| = r_k\}$, where x is a k-seed. Let W_k be the union of all k-cutters, and define a ***k*-blob** to be any connected component of $\mathbb{R}^d \setminus \bigcup_{j \geq k} W_j$. Clearly, every M-point lies in exactly one k-blob for each k almost surely, and every k-blob is a subset of exactly one $(k+1)$-blob. (Note that k-cutters typically occur in nearly-coincident pairs. This fact makes the pictures rather odd, but neither helps nor hinders our proof. The anomaly could be avoided at the expense of a less convenient definition of k-seeds).
For each k, define a partition $\mathcal{P}_k$ of $[M]$ by declaring two M-points to be in the same clump of $\mathcal{P}_k$ if they lie in the same k-blob. Clearly $\mathcal{P}_1, \mathcal{P}_2, \ldots$ form a clumping; we must check that it is locally finite and connected.
First we claim that almost surely all blobs are bounded. This will imply immediately that the clumping is locally finite. It is sufficient to check that almost surely all blobs which intersect $B(1)$ are bounded. Let

$$V_k = \{B(1) \text{ is enclosed by some } k\text{-cutter}\}.$$

We will show that $\mathbf{P}[V_k] \to 1$ as $k \to \infty$. This implies that V_k occurs for infinitely many k almost surely, and this implies the claim, since any j-blob which intersects $B(1)$ must then be enclosed by some k-cutter for some $k \geq j$. By the definition of a k-cutter, V_k equals the event that $B(r_k - 1)$ contains some k-seed. Now, we have

$$\begin{aligned}\mathbf{P}\left[B(a_k/2) \text{ contains some } k\text{-seed}\right] &\geq \mathbf{P}\left[M(B(a_k/2)) = 2\right] \\ &= e^{-C_2 a_k^d}(C_2 a_k^d)^2/2 \geq C_3 a_k^{2d}\end{aligned}$$

Moreover, if $A, B \subseteq \mathbb{R}^d$ are sets at distance at least $2a_k$ from each other, then the events $\{A \text{ contains some } k\text{-seed}\}, \{B \text{ contains some } k\text{-seed}\}$ are independent. For $k \geq 2$ we may clearly find $\lceil C_4(r_k/a_k)^d \rceil$ balls of radius $a_k/2$ lying in $B(r_k - 1)$ and spaced at distance at least $2a_k$ from each other. Hence

$$\begin{aligned}\mathbf{P}\left[V_k\right] = \mathbf{P}\left[B(r_k - 1) \text{ contains some } k\text{-seed}\right] &\geq 1 - (1 - C_3 a_k^{2d})^{C_4(r_k/a_k)^d} \\ &\geq 1 - e^{-C_5 a_k^d r_k^d} = 1 - e^{-C_5 e^{k/2}} \to 1,\end{aligned}$$

establishing the above claim.
To prove that the clumping is connected, it suffices to prove that for every fixed $\ell > 0$, almost surely all M-points in $B(\ell)$ lie in the same clump of $\mathcal{P}_k$ for some k. By the construction it is enough to show that all such M-points lie in a single blob, and this in turn follows if almost surely all of $B(\ell)$ lies in some blob. Let

$$U_k = \{B(\ell) \text{ intersects some } k\text{-cutter}\}.$$

We have

$$\begin{aligned}\sum_{k=1}^{\infty} \mathbf{P}\left[U_k\right] &= \sum_{k=1}^{\infty} \mathbf{P}\left[B(r_k + \ell) \setminus B(r_k - \ell) \text{ contains some } k\text{-seed}\right] \\ &\leq \sum_{k=1}^{\infty} \lambda_k C_6 \ell r_k^{d-1} \leq C_7 \ell \sum_{k=1}^{\infty} e^{-k/2} \qquad \text{(by (1))} \\ &< \infty,\end{aligned}$$

so the Borel-Cantelli lemma implies that $\mathbf{P}\left[U_k \text{ occurs for infinitely many } k\right] = 0$, and hence $B(\ell)$ lies in some blob as required. ☐

3 Forests and Matchings

Lemma 8 *Let $d \geq 1$ and let M be a non-equidistant point process in $\mathbb{R}^d$ which is invariant and ergodic under isometries. The following are equivalent.*

(i) M has a clumping which almost surely is locally finite and has each component infinite.

(ii) M has a factor graph in which almost surely each component is a locally finite one-ended tree.

(iii) M has a factor graph in which almost surely each component is a directed doubly infinite path.

(iv) M has a factor graph in which almost surely each component is a locally finite one-ended or two-ended tree.

PROOF. The implications (i)⇒(ii) and (ii)⇒(iii) follow immediately by applying the proofs of Lemmas 6,7 to each component. The implication (iii)⇒(iv) is a triviality since a doubly infinite path is a two-ended tree. Therefore it is sufficient to prove (iv)⇒(i).
Suppose G is a factor graph as in (iv). We will treat each component separately. Let C be a component of G. If C is a one-ended tree, let π be any singly infinite path in C, and define

$$L(C) = \limsup\{\|x-y\| : (x,y) \text{ is an edge of } \pi\}.$$

Since any two singly infinite paths in C must eventually coalesce, $L(C)$ does not depend on the choice of π. On the other hand if C is a two-ended tree, then it has a unique doubly infinite self-avoiding path, which we call the **trunk**. If the trunk is deleted, only finite components remain. Define

$$L(C) = \sup\{\|x-y\| : (x,y) \text{ is an edge of the trunk}\}.$$

Let k be a positive integer. Define for each component C

$$L_k(C) = \begin{cases} L(C) - k^{-1} & \text{if } L(C) < \infty, \\ k & \text{if } L(C) = \infty. \end{cases}$$

Let G_k be the graph obtained from G by deleting every edge (x,y) for which $\|x-y\| \in [L_k(C), L(C))$, where C is the component containing (x,y). Let $\mathcal{P}_k$ be the partition of $[M]$ induced by the components of G_k. Clearly $\mathcal{P}_1, \mathcal{P}_2, \ldots$ form a clumping; we claim that it is locally finite and has the same components as G; this will establish (i).
First note that there is no edge (x,y) for which $\|x-y\| = L(C)$ where C is the component of G containing (x,y). To see this, note that by non-equidistance, there can be at most one such edge in each component of G. But now consider the mass transport in which every vertex in component C sends one unit of mass to each of x,y if there is such an edge (x,y) in C. Applying Lemma 5 gives a contradiction since each vertex sends out at most two units, but such vertices x,y would receive infinite mass.
It follows that if x,y are vertices in the same component C then there exists some k such that no edge in the path from x to y has length lying in $[L_k(C), L(C))$, so x,y lie in the same clump of $\mathcal{P}_k$. Hence the components of the clumping are the components of G.
It remains to show that the all components of G_k are finite, since this will imply that the clumping is locally finite. If C is a one-ended component of G, the definition of $L(C)$ implies that every infinite path in G has some edge with length in $[L_k(C), L(C))$, so all components of G_k which lie in one-ended components of G are finite. Let H be the graph consisting of all trunks of two-ended components of G, and let $H_k = G_k \cap H$. Clearly all components of H_k are doubly infinite paths, singly infinite paths, or finite paths. We claim that in fact the first two possibilities can be ruled out, and this implies that all components of G_k must be finite. We prove the claim as follows. Firstly, the definition of $L(C)$ for a two-ended component C of G implies that the trunk of C has some edge with length in $[L_k(C), L(C))$, so components of H_k cannot be doubly infinite paths. Secondly, consider the mass transport in which every vertex in a singly infinite path component of H_k sends one unit of mass to the end point of the singly infinite path. Each vertex sends out at most one unit, but if singly infinite paths existed then their end points would receive infinite mass, violating Lemma 5. □

PROOF OF THEOREM 2. The **minimal spanning forest** S of M is the factor graph obtained from the complete graph on $[M]$ by deleting every edge which is the longest in some cycle. It is proved in [4] that all components of S are one-ended or two-ended trees almost surely. And S is locally finite by Lemma 4. Hence Theorem 2 (i),(ii) follow from Lemma 8.

Let $\mathcal{P}_1, \mathcal{P}_2, \ldots$ be a clumping as in Lemma 8 (i). We will construct a perfect matching, establishing Theorem 2 (iii). First, consider a clump $A \in \mathcal{P}_1$ which has at least two M-points, match the two closest M-points of A, then remove them and repeat until A has at most one unmatched M-point. Do this for every clump $A \in \mathcal{P}_1$. Next apply the same construction to each clump $A \in \mathcal{P}_2$ which has at least two unmatched M-points, and so on indefinitely. It is clear that each component of the clumping contains at most one M-point that is never matched. We claim that in fact there are no such M-points. To prove this consider the mass transport in which x sends one unit to y if x, y are in the same component of the clumping and y is never matched. Since all components of the clumping are infinite, any M-point never matched would receive infinite mass, contradicting Lemma 5. □

Extension to more general point processes. As mentioned in the introduction, the condition of non-equidistance in Theorem 2 may be relaxed to the condition that the symmetry group of $[M]$ is almost surely trivial. The proof above goes through once we have defined a real-valued measurable function $\varphi(x, \mu)$ where $x \in \mathbb{R}^d$ and μ is a realization of M, such that $\varphi(\theta x, \theta\mu) = \varphi(x, \mu)$ for any isometry θ of $\mathbb{R}^d$, and $\varphi(x, \mu) \neq \varphi(y, \mu)$ for $x \neq y$. This function is used to break ties between equidistant pairs: if $\|x - y\| = \|x - w\|$ then the pair $\{x, y\}$ is preferred to $\{x, w\}$ provided that $\varphi(y, \mu) < \varphi(w, \mu)$. Similarly, if $\|x - y\| = \|z - w\|$ where these points are distinct, then $\{x, y\}$ is preferred if and only if $\min\{\varphi(x, \mu), \varphi(y, \mu)\} < \min\{\varphi(z, \mu), \varphi(w, \mu)\}$.

It remains to define φ. An **open dyadic cube** in $\mathbb{R}^d$ is a set of the form $2^t[z + (0,1)^d]$ for $t \in \mathbb{Z}$ and $z \in \mathbb{Z}^d$. Let $\{Q_n\}_{n\geq 1}$ be an enumeration of all open dyadic cubes (of all scales) in $\mathbb{R}^d$. Fix μ. Denote by $O(d)$ the collection of isometries preserving the origin, that is, orthogonal matrices. For $\rho \in O(d)$, define $\psi_n(\rho, x) = 1$ if $\mu(x + \rho Q_n) \geq 1$ and $\psi_n(\rho, x) = 0$ otherwise. Each ψ_n is lower semi-continuous, and hence so is the function

$$\psi(\rho, x) = \sum_{n\geq 1} 3^{-n}\psi_n(\rho, x)\,.$$

The assumption on the symmetry group of μ implies that almost surely, for every $x \neq y$ in $\mathbb{R}^d$ and any pair $\rho_1, \rho_2 \in O(d)$, there is an n such that $\psi_n(\rho_1, x) \neq \psi_n(\rho_2, y)$, whence $\psi(\rho_1, x) \neq \psi_n(\rho_2, y)$. Finally, define

$$\varphi(x, \mu) = \min_{\rho\in O(d)} \psi(\rho, x)\,.$$

4 Iterated Nearest Neighbor Matching

In this section we prove Theorem 3. The mass transport principle in Lemma 5 extends to the more general setting, with any bounded Borel set K in place of the unit cube. In the case of the hyperbolic plane the result is a special case of Theorem 5.2 in [6], and the proof in [6] extends to our more general setting. Lemma 8 extends to the general setting with essentially the same proof. In the case when Γ is amenable, the proof of Theorem 2 also extends, since all the components of any invariant random forest must have either one or two ends. In the non-amenable case, this argument breaks down; although it is believed that the components of the (wired) minimal spanning forest all have at most two ends, this has not been proved. Thus we use a different construction, which is of interest even in Euclidean space.

Let ρ be the metric on Λ. By a **descending chain** we mean a sequence of distinct M-points $x_1, x_2, \ldots$ for which the distances $\rho(x_i, x_{i+1})$ form a (strictly) decreasing sequence. We will see that the argument becomes simpler in the case when M has no descending chains, and in particular we obtain a perfect matching by the natural procedure of iteratively matching mutually nearest neighbors. In particular it was proved in [10] that Poisson processes have no descending chains almost surely. On the other hand, there do exist isometry-invariant point processes with descending chains. For example, start with a Poisson process on $\mathbb{R}^d$, construct a one-ended tree as a factor graph (Theorem 1), and then add extra points along the edges in such a way that every singly infinite path becomes a descending chain.

Stable matchings. A perfect matching of $[M]$ is called **stable** if for every pair of edges $(x,x'),(y,y')$ we have $\rho(x,y) \geq \min\{\rho(x,x'),\rho(y,y')\}$. This definition was made, in a more general context, by Gale and Shapley [8]. When M has no descending chains, we will see that iteratively matching mutually nearest neighbors yields the unique stable perfect matching.

PROOF OF THEOREM 3. We start by proving that M has a perfect matching as a factor graph. Then we will deduce the statements (i),(ii) in Theorem 2.
Call a pair of M-points x,y **mutually closest** if x is the closest M-point to y and y is the closest M-point to x. We may construct a (not necessarily perfect) matching G as a factor graph of M by the following procedure. First match all mutually closest pairs of M-points, and remove such points, then match and remove all mutually closest pairs among the remaining points, and repeat indefinitely. Let N be the process of all M-points which are never matched by this procedure (that is, that have degree 0 in G). Clearly N is an isometry-invariant ergodic point process, so in particular it has almost surely infinitely many points or almost surely no points. The matching G is perfect if and only if N has no points. (We shall see below that this indeed holds if M has no descending chains. The above procedure was suggested by Dana Randall, and is noted in [7]. It has the following informal interpretation. Imagine a growing ball centered at each M-point, such that at time t each ball has radius t. Every time two balls meet, they are annihilated and their centers are matched. From the iterative description, one may verify by induction that every edge of G is present in every stable perfect matching).
Consider the directed factor graph H of N in which there is a directed edge from each N-point to its closest N-point. It is easy to see that H has no cycles except of size 2 (these being exactly the mutually closest pairs of N-points), and that every finite component of H contains exactly one cycle of size 2. We claim that in fact N has no mutually closest pairs, so H has no finite components. To see this, suppose that x,y are mutually closest N-points. This is equivalent to the statement that the set $J = \{z \in \Lambda : \rho(z,x) \leq \rho(x,y) \text{ or } \rho(z,y) \leq \rho(x,y)\}$ contains no N-points other than x,y. Since J is bounded, it contains only finitely many M-points almost surely, and hence at some finite (random) stage of the matching procedure above, J contains no unmatched M-points other than x,y. But then x,y will be matched at the next stage, which contradicts the assumption that they are N-points.
The above argument shows that if G is not a perfect matching then H is non-empty and every component of H is an infinite tree. (Note in particular that the latter is possible only if M has descending chains). We claim that the components of H are one-ended or two-ended trees. Once this is established we obtain a perfect matching of M by using Lemma 8 (extended to Λ) to perfectly match N, and then combining this with G.
To prove the above claim, define the **backbone** B of H to be the directed subgraph of N whose edge set is the union of all (not necessarily directed) doubly infinite self-avoiding paths of H. Note that all vertices incident to B have degree at least 2 in B. The claim is equivalent to the assertion that B has no vertices of degree greater than 2. Consider the mass transport in which x sends one unit of mass to y if B has a directed edge from x to y. Note that the

degree of a vertex in B equals the total mass sent out plus the total mass received. By the construction of H, each vertex sends out at most one unit. Let K be a fixed bounded Borel set, let V be the set of N-points in K incident to B, and let $D = \sum_{x \in V} \deg_B(x)$. Then the mass transport principle yields

$$\mathbf{E}\, D \leq 2\mathbf{E}\,|V|.$$

On the other hand, if B has vertices of degree greater than 2 then

$$\mathbf{E}\,[D - 2|V|] > 0,$$

a contradiction, proving the claim.

The above argument shows that every M satisfying the conditions of Theorem 3 has a perfect matching as a factor graph. We deduce that statements (i),(ii) of Theorem 2 hold as follows. Construct a perfect matching, then delete one M-point from each matched pair (choosing which one as in the choice of leaders in the proof of Lemma 6, for example), then construct a perfect matching as a factor graph of the process of remaining points, and repeat indefinitely. This gives a locally finite clumping with all components infinite, so we can apply Lemma 8 (extended to Λ). □

5 Open Problems

(i) Consider a perfect matching which is a factor graph of a point process. What can be said about the lengths of the edges? More specifically, how does the probability that some point in $B(1)$ is matched to a point outside $B(r)$ behave as $r \to \infty$? One may ask such questions for specific matchings such as those discussed in Section 4, or one may ask for the optimal tail behavior over all possible perfect matchings. The latter question is related to problems studied in [12], for example.

(ii) When does the iterated matching construction in the last paragraph of the proof of Theorem 3 result in a single one-ended tree, as opposed to a forest? In particular, for which d does the Poisson process on $\mathbb{R}^d$ yield a tree? (The case $d = 1$ is easy).

(iii) Steve Evans suggested the following questions. What other graphs are possible as factor graphs of point processes? For example, for which d, n does the Poisson process on $\mathbb{R}^d$ have $\mathbb{Z}^n$ as a factor graph?

(iv) For what other point processes do the conclusions of Theorem 1 hold? In particular, do they hold for every non-equidistant ergodic point process on $\mathbb{R}^d$?

Remark added in revision: Russ Lyons has informed us that Adam Timar has made progress on problems (iii),(iv) above.

Acknowledgments

We thank David Aldous, Dana Randall and Hermann Thorisson for valuable discussions, and the anonymous referee for helpful comments.

References

[1] S. Adams and R. Lyons. Amenability, Kazhdan's property and percolation for trees, groups and equivalence relations. *Israel J. Math.*, 75(2-3):341–370, 1991.

[2] S. R. Adams and R. J. Spatzier. Kazhdan groups, cocycles and trees. *Amer. J. Math.*, 112(2):271–287, 1990.

[3] D. Aldous and J. M. Steele. Asymptotics for Euclidean minimal spanning trees on random points. *Probab. Theory Related Fields*, 92(2):247–258, 1992.

[4] K. S. Alexander. Percolation and minimal spanning forests in infinite graphs. *Ann. Probab.*, 23(1):87–104, 1995.

[5] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.

[6] I. Benjamini and O. Schramm. Percolation in the hyperbolic plane. *J. Amer. Math. Soc.*, 14(2):487–507 (electronic), 2001.

[7] P. A. Ferrari, C. Landim, and H. Thorisson. Poisson trees, succession lines and coalescing random walks. Preprint.

[8] D. Gale and L. Shapley. College admissions and stability of marriage. *Amer. Math. Monthly*, 69(1):9–15, 1962.

[9] O. Häggström. Infinite clusters in dependent automorphism invariant percolation on trees. *Ann. Probab.*, 25(3):1423–1436, 1997.

[10] O. Häggström and R. Meester. Nearest neighbor and hard sphere models in continuum percolation. *Random Structures Algorithms*, 9(3):295–315, 1996.

[11] R. Pemantle and Y. Peres. Nonamenable products are not treeable. *Israel J. Math.*, 118:147–155, 2000.

[12] J. M. Steele. *Probability theory and combinatorial optimization*, volume 69 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.

[13] H. Thorisson. *Coupling, stationarity, and regeneration*. Probability and its Applications. Springer-Verlag, New York, 2000.