The Annals of Probability
2017, Vol. 45, No. 1, 469–517
DOI: 10.1214/14-AOP935

FROM RANDOM LINES TO METRIC SPACES

BY WILFRID S. KENDALL[1]

University of Warwick

This paper is dedicated to the memory of Don Burkholder, a great probabilist and a kind man

Consider an improper Poisson line process, marked by positive speeds so as to satisfy a scale-invariance property (actually, scale-equivariance). The line process can be characterized by its intensity measure, which belongs to a one-parameter family if scale and Euclidean invariance are required. This paper investigates a proposal by Aldous, namely that the line process could be used to produce a scale-invariant random spatial network (SIRSN) by means of connecting up points using paths which follow segments from the line process at the stipulated speeds. It is shown that this does indeed produce a scale-invariant network, under suitable conditions on the parameter; in fact, it then produces a parameter-dependent random geodesic metric for d-dimensional space ($d \geq 2$), where geodesics are given by minimum-time paths. Moreover, in the planar case, it is shown that the resulting geodesic metric space has an almost everywhere unique-geodesic property that geodesics are locally of finite mean length, and that if an independent Poisson point process is connected up by such geodesics then the resulting network places finite length in each compact region. It is an open question whether the result is a SIRSN (in Aldous' sense; so placing finite *mean* length in each compact region), but it may be called a pre-SIRSN.

1. Introduction. Recent work in random spatial networks [2, 3] has focussed on specification and analysis of an intriguing class of random networks known as *scale-invariant random spatial networks* (SIRSN). Motivated by the success of Google Maps and Bing Maps, Aldous [2] shows how a natural collection of desirable properties (statistical invariance under translation, rotation and scale-change and some integrability conditions) define a class of models with a useful structure theory.

DEFINITION 1.1 (Definition of a SIRSN, Aldous [2]). Consider a d-dimensional random mechanism, which provides random routes connecting any

Received March 2014; revised April 2014.
[1]Supported in part by EPSRC Research Grant EP/K013939.
MSC2010 subject classifications. Primary 60D05; secondary 90B15, 90B20, 46E35.
Key words and phrases. Fiber process, line-space, Lipschitz path, marked line process, orientation field, perpetuity, Poisson line process, pre-SIRSN, random upper-semicontinuous function, Π-geodesic, Π-path, scale invariance, SIRSN, Sobolev space, spatial network, stochastic geometry, weak SIRSN.

two points $x_1, x_2 \in \mathbb{R}^d$. We say that this is a *SIRSN* if the following properties hold:

1. Between any specified two points $x_1, x_2 \in \mathbb{R}^d$, almost surely the random mechanism provides just one connecting random route $\mathcal{R}(x_1, x_2) = \mathcal{R}(x_2, x_1)$, which is a finite-length path connecting x_1 to x_2.
2. For a finite set of points $x_1, \ldots, x_k \in \mathbb{R}^d$, consider the random network $\mathcal{N}(x_1, \ldots, x_k)$ formed by the random routes provided by the structure to connect all x_i and x_j. Then $\mathcal{N}(x_1, \ldots, x_k)$ is statistically invariant (strictly speaking, equivariant) under translation, rotation, and rescaling: if $\mathcal{S}$ is a Euclidean similarity of $\mathbb{R}^d$ then the networks $\mathcal{S}\mathcal{N}(x_1, \ldots, x_k)$ and $\mathcal{N}(\mathcal{S}x_1, \ldots, \mathcal{S}x_k)$ have the same distribution.
3. Let D_1 be the length of the route between two points separated by unit Euclidean distance. Then $\mathbb{E}[D_1] < \infty$.
4. Suppose that Ξ_λ is a Poisson point process in $\mathbb{R}^d$, of intensity $\lambda > 0$ and independent of the random mechanism in question. Then $\mathcal{N}(\Xi_\lambda)$, the union of all the networks $\mathcal{N}(x_1, \ldots, x_k)$ for $x_1, \ldots, x_k \in \Xi_\lambda$, is a locally finite fiber process in $\mathbb{R}^d$. That is to say, for any compact set K the total length of $\mathcal{N}(\Xi_\lambda) \cap K$ is almost surely finite.
5. The length intensity ℓ of $\mathcal{N}(\Xi_1)$ (the mean length per unit area) is finite.
6. Suppose the Poisson point processes $\{\Xi_\lambda : \lambda > 0\}$ are coupled so that $\Xi_{\lambda_1} \subseteq \Xi_{\lambda_2}$ if $\lambda_1 < \lambda_2$. The fiber process

$$\bigcup_{\lambda>0} \bigcup_{x_1,x_2\in\Xi_\lambda} \left(\mathcal{R}(x_1, x_2) \setminus \left(\mathrm{ball}(x_1, 1) \cup \mathrm{ball}(x_2, 1)\right)\right)$$

has length intensity bounded above by a finite constant $p(1)$.

If only properties 1–5 are satisfied, then the random mechanism is called a *weak SIRSN*. If only properties 1–4 are satisfied, then the random mechanism is called a *pre-SIRSN*.

Aldous and Ganesan [3] describe the *binary hierarchy model*, a structure for providing planar routes, based on minimum-time paths using a dyadic grid furnished with speeds and uniformly randomized in orientation and position. Aldous [2] proves that this is a full planar SIRSN satisfying all the requirements of Definition 1.1. Aldous and Ganesan [3] also propose two other candidates for planar SIRSNs which do not involve the somewhat unnatural randomization required for the binary hierarchy model: one is based on route-provision via a scale-invariant improper Poisson line process marked with random speeds (the *Poisson line process model*); and the other uses a dynamic proximity graph related to the Gabriel graph. The purpose of the present paper is to explore the Poisson line process model: we will show that it is at least a pre-SIRSN if $d = 2$, and moreover, we will show that even in dimension $d > 2$ the construction provides a random metric space on $\mathbb{R}^d$ (in particular it satisfies at least properties 1 and 2 of Definition 1.1,

with the possible exception of uniqueness of route). This therefore establishes the significance of the Poisson line process model as a scale-invariant random spatial network, while leaving open the question of whether it is a weak SIRSN or even a full SIRSN, not just a pre-SIRSN.

The chief difficulty in analyzing any of these random mechanisms lies in the fact that it is hard to work with explicit minimum-time paths, whose explicit construction would involve solving a nonlocal minimization problem to determine geodesics. Aldous and Kendall [5] and Kendall [17, 19] use approximations known as "near-geodesics," constructed using a kind of greedy algorithm. Baccelli, Tchoumatchenko and Zuyev [7] and Broutin, Devillers and Hemsley [11] study Delaunay tessellation paths that are determined using either their relationship to appropriate Euclidean straight lines or the so-called "cone walk." LaGatta [20] studies geodesics determined by random *smooth* Riemannian structures, for which conventional calculus methods are available. In the following, we argue for existence of minimum-time paths by exploiting properties of a Sobolev space of paths, and then by using indirect arguments.

The structure of the paper is as follows. The rest of this Introduction (Section 1) is concerned with basic notions of stochastic geometry (Section 1.1) and with the definition of the underlying improper Poisson line process Π marked with speeds (Section 1.2). This improper Poisson line process Π is defined by an intensity measure $(\gamma - 1)v^{-\gamma}\,\mathrm{d}v\mu_d(\mathrm{d}\ell)$ (for speed $v > 0$, parameter $\gamma > 1$, and invariant measure μ_d on line-space) and supplies a measurable orientation field marked by speeds: Section 2 then explores the way in which the measurable orientation field can be integrated to provide Lipschitz paths based on the marked line process, namely Π-paths. Sobolev space and comparison arguments can then be used to establish a priori bounds on Lipschitz constants for finite-time Π-paths (Theorem 2.6), hence closure, weak closure and finally weak compactness (Corollary 2.11) of finite-time Π-paths. All these results require $\gamma \geq d$. Note that dimension $d > 1$ is required if line-process theory is to be nontrivial.

Section 3 shows that, given $\gamma > d$ and fixed points $x_1, x_2 \in \mathbb{R}^d$, it is almost surely possible to connect x_1 to x_2 in finite time with Π-paths (Theorem 3.1), and indeed with probability 1 it is possible to connect *all* pairs of points in this way (Theorem 3.6). Combined with Corollary 2.11, this implies the existence of minimum-time Π-paths, namely Π-*geodesics* (Definition 3.4, Corollary 3.5). In dimension $d > 2$, this is a rather unexpected result, since almost surely none of the lines of Π will then intersect. Nevertheless, Π then furnishes $\mathbb{R}^d$ with the structure of a random geodesic metric space. In these higher dimensions it is difficult to imagine what a Π-geodesic might look like (Figure 1 illustrates the easier $d = 2$ case): however Definition 3.7, Theorem 3.8 and Corollary 3.9 describe a class of "ε-near-sequential-Π-paths" which can be used to approximate (and to simulate) Π-geodesics (Theorem 3.11). In particular, these results imply measurability of the random time taken to pass from one point to another using a Π-geodesic (Corollary 3.12).

The remainder of the paper is restricted to the planar case of $d = 2$, since the arguments now make essential use of point-line duality. Consider the extent to which networks formed by Π-geodesics fulfil the requirements of Definition 1.1. The statistical invariance property 2 follows immediately from similar invariance of the underlying intensity measure of the improper Poisson line process (whether planar or not). Property 1 requires almost sure uniqueness of network routes: Section 4 establishes this for $\gamma > d = 2$ (Theorem 4.4), using a careful analysis of the nature of planar Π-geodesics (Theorem 4.3) which falls just short of establishing that planar Π-geodesics can be made up of consecutive sequences of line segments. While Π-geodesics between pairs of points are minimum-time paths, the fact that they have finite mean length is not immediately apparent; this is established in Section 5, first for restricted planar Π-geodesics (Lemma 5.1), then for general planar Π-geodesics (Theorem 5.2). Thus the finite-mean-length property 3 of Definition 1.1 is verified for $d = 2$. Finally the pre-SIRSN property 6 is established for the planar case in Theorem 6.4 of Section 6; here also is established the uniqueness of planar Π-geodesics reaching out to infinity (Theorem 6.2) and, for any specified point $x \in \mathbb{R}^2$, the fact that all Π-geodesics emanating from x must coincide for initial periods (Theorem 6.3). These results are established using an essentially soft argument concerning the existence of certain structures in Π (Lemma 6.1); the concluding Section 7 notes that more quantitative arguments would be required to decide whether the weak SIRSN or full SIRSN properties hold—indeed Jonas Kahn has very recently shown me such arguments. Section 7 also notes some other interesting open questions.

1.1. *Notation and basic results for random line processes.* Random line processes (random patterns of lines) play a fundamental role in this study. Here, we review notation and basic results for unsensed random line processes in Euclidean space, as described in Chiu et al. [13], Chapter 8. (By an "unsensed line," we mean a line without preferred direction.) The corresponding theory for sensed lines follows from the observation that the space of sensed lines forms a double cover of the space of unsensed lines.

Consider *line-space*, the space $\mathcal{L}^d$ of all unsensed lines in $\mathbb{R}^d$, for dimension $d \geq 2$. In the planar case $d = 2$, there is a natural geometric representation of $\mathcal{L}^2$ as a punctured projective plane, since there is a 3-space construction of the family of planar lines as the family of intersections of 2-subspaces with a reference plane (say $x_3 = 1$). More visually, but less naturally, $\mathcal{L}^2$ can be viewed as a Möbius band of infinite width. Similar but less graphic geometric descriptions of $\mathcal{L}^d$ (and its sensed counterpart) can be given in higher dimensional cases ($d > 2$): for example, the space of *sensed* lines in $\mathbb{R}^d$ can be represented using the standard immersion of the tangent bundle $T S^{d-1}$ of the $(d-1)$-sphere in $\mathbb{R}^d$.

It is convenient to introduce notation for hitting events and hitting sets. For a line $\ell \in \mathcal{L}^d$ and for K, a compact subset of $\mathbb{R}^d$, we write

$$\ell \Uparrow K \tag{1.1}$$

for the statement that ℓ intersects K. We also introduce the *hitting set* of K (the set of lines that hit K):

$$[K] = \{\ell \in \mathcal{L}^d : \ell \Uparrow K\}. \tag{1.2}$$

General arguments show that there exists a measure on $\mathcal{L}^d$ that is invariant under Euclidean isometries and unique up to a scaling factor. Line-space $\mathcal{L}^d$ can be constructed as the quotient space of the group of d-dimensional rigid motions by the subgroup that leaves invariant a specified line. The existence of invariant measure on line-space follows from the study of quotient measures for locally compact topological groups; a conceptual and general treatment of existence and uniqueness is given by Abbaspour and Moskowitz [1], Section 2.3 (see also Loomis [21], pages 130–133), and follows here from unimodularity of the two groups in question. Santaló [23], Chapter 10, and Ambartzumian [6] describe alternative approaches that are direct but are computational rather than conceptual.

DEFINITION 1.2. *Invariant line measure* $\mu_d(\mathrm{d}\ell)$ is the unique measure on $\mathcal{L}^d$ that is invariant under Euclidean isometries and is normalized by the following requirement: for all compact convex sets $K \subset \mathbb{R}^d$ of nonempty interior ("convex bodies"), the μ_d-measure of the hitting set $[K]$ is half the Hausdorff $(d-1)$-dimensional measure of the boundary of K:

$$\mu_d\big([K]\big) = \tfrac{1}{2} m_{d-1}(\partial K). \tag{1.3}$$

Here and in the following, m_{d-1} denotes Hausdorff measure of dimension $d-1$. The purpose of the normalization factor $\frac{1}{2}$ is to ensure that the μ_d-measure of the hitting set of a fragment A of a flat hypersurface is equal to its hypersurface area $m_{d-1}(A)$.

In the important special case of $d = 2$, we can parameterize an unsensed line $\ell \in \mathcal{L}^2$ by (a) the angle $\theta = \theta(\ell) \in [0, \pi)$ that it makes with a reference line (say, the x-axis), and (b) the *signed* distance $r = r(\ell)$ between the line ℓ and a reference point [conventionally taken to belong to the reference line; say, the origin $\mathbf{o} = (0, 0)$]. Equation (1.3) then takes a more explicit form:

$$\mu_2(\mathrm{d}\ell) = \tfrac{1}{2}\,\mathrm{d}r\,\mathrm{d}\theta. \tag{1.4}$$

More generally, the line measure $\mu_d(\mathrm{d}\ell)$ can be disintegrated using $(d-1)$-dimensional Hausdorff measure on the hyperplane perpendicular to ℓ. Let ϖ be the unsensed direction of ℓ and let y be its point of intersection on the perpendicular hyperplane. Let $\kappa_s = \frac{s/2}{\Gamma(1+s/2)}$ denote the s-dimensional volume of the unit ball in $\mathbb{R}^s$, and for later convenience let $\omega_{s-1} = s\kappa_s$ denote the hypersurface area of its boundary. Then

$$\mu_d(\mathrm{d}\ell) = \frac{1}{\kappa_{d-1}} m_{d-1}(\mathrm{d}y) m_{S_+^{d-1}}(\mathrm{d}\varpi), \tag{1.5}$$

where the measure $m_{S_+^{d-1}}$ is defined on the space of unsensed line directions and can be thought of as $(d-1)$-dimensional Hausdorff measure on the unit hemisphere S_+^{d-1} in $\mathbb{R}^d$. Proper interpretation of the representation (1.5) requires the space of unsensed directions to be considered as a further projective space, and the product measure to be twisted to take account of the fact that m_{d-1} here is defined on the hyperplane normal to the unsensed direction of the line in question. However, the resulting discrepancies are confined to a null-set which can be ignored when considering invariant Poisson line processes.

An alternative representation, useful for certain calculations, describes μ_d in terms of the intersection of ℓ with a fixed reference hyperplane. In two dimensions, we obtain

$$\mu_2(\mathrm{d}\ell) = \tfrac{1}{2}\sin\theta\,\mathrm{d}p\,\mathrm{d}\theta, \tag{1.6}$$

where $p = p(\ell)$ is the signed distance from the reference point $\mathbf{o}$ to the intersection of ℓ with the reference line. This alternative representation is defective: if $\theta = 0$ then there is no intersection and so p is ill-defined. However once again the resulting discrepancies are confined to a null-set which can be ignored when considering invariant Poisson line processes. In higher dimensions, the corresponding representation is

$$\mu_d(\mathrm{d}\ell) = \frac{\sin\theta}{\kappa_{d-1}} m_{d-1}(\mathrm{d}z) m_{S_+^{d-1}}(\mathrm{d}\varpi), \tag{1.7}$$

where θ is the angle made by the unsensed direction ϖ of the line ℓ with the fixed reference hyperplane, and z locates the intersection of ℓ with the reference hyperplane.

Note finally that the arguments of this paper depend only on the general forms of equations (1.4)–(1.7); the exact multiplicative constants involved are not crucial.

1.2. *Improper Poisson line processes.* Our constructions use *Poisson* line processes. A unit-intensity Poisson line process in $\mathbb{R}^d$ is obtained simply by generating a Poisson point process on the corresponding representing space $\mathcal{L}^d$ using the invariant measure μ_d. It is a geometric consequence of the σ-finiteness of μ_d that the resulting random line pattern is locally finite: only finitely many lines hit any given compact set. However, our constructions will use *improper* Poisson line processes, which can be viewed as superpositions of infinitely many independent Poisson line processes, different line processes being thought of as representing highways with speed limits lying in different ranges. If we augment the representation space by a mark space $(0,\infty)$ of speed limits, then the improper Poisson line process can be represented as a Poisson point process on $\mathcal{L}^d \times (0,\infty)$, with a σ-finite intensity measure on $\mathcal{L}^d \times (0,\infty)$ which is invariant under rigid motions but which does not project down onto a σ-finite intensity measure on $\mathcal{L}^d$. Thus, the main actors in this account are invariant improper unsensed Poisson line processes,

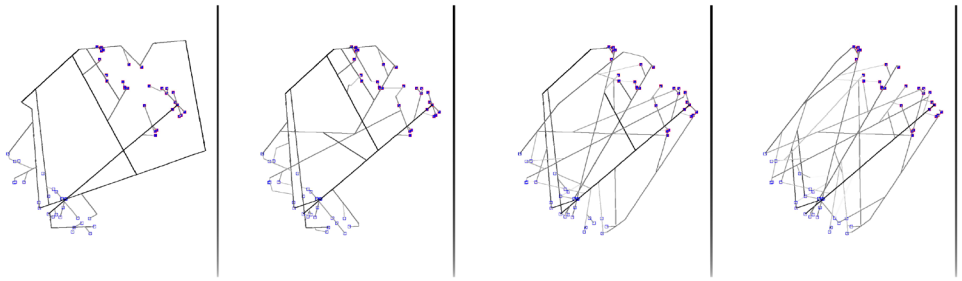

FIG. 1. *Minimum-time routes between two separated collections of nodes for networks built from the improper Poisson line process described in Section 1.2, for parameter γ taking values 2.1, 4.0, 8.0, 16.0. Lighter segments have lower speed limits. Note that as γ increases so the routes become more direct, but also there is less route-sharing.*

with each line ℓ being marked by a different positive speed limit $v = v(\ell) > 0$. Scaling arguments [2, 3] lead to a natural family of intensity measures for such a marked line process, based on a positive parameter $\gamma > 1$:

$$(\gamma - 1)v^{-\gamma}\,\mathrm{d}v\mu_d(\mathrm{d}\ell). \tag{1.8}$$

The factor $\gamma - 1$ ensures that for all $\gamma > 1$ the subprocess of lines with marks $v > 1$ forms a unit-intensity Poisson line process which is of *unit intensity*, in the sense that its mean intensity is the invariant measure given in (1.3), so that the mean number of lines hitting a flat fragment of hypersurface is equal to its hypersurface area. In case $d = 2$, we may write this intensity measure as $\frac{1}{2}(\gamma - 1)v^{-\gamma}\,\mathrm{d}v\,\mathrm{d}r\,\mathrm{d}\theta$. Fixing a general dimension d and parameter $\gamma > 1$, let $\Pi = \Pi^{(d,\gamma)}$ denote the resulting random process of marked lines $(\ell, v(\ell))$. In the following, the dependence on d and γ will be clear from the context, and consequently will be suppressed. Figure 1 illustrates the formation of minimum-time routes between two fixed collections of nodes, for varying values of the parameter $\gamma > 2$. Note that spatial networks formed in this way will automatically satisfy property 2 of Definition 1.1, because of the invariance properties of the intensity measure (1.8).

The intensity measure gives infinite mass to the set of lines intersecting any convex body and, therefore, the union of all lines from Π is *not* a random closed set. Consequently, it is not possible to make direct application of the classic theory of random closed sets (as surveyed, e.g., in Chiu et al. [13], Chapter 6). Indeed the union of all lines from Π is almost surely everywhere dense, and the theory for such sets is obscure (see, e.g., Aldous and Barlow [4], Kendall [18]). We therefore focus on subpatterns of lines subject to a positive lower bound on their speeds. Consider the set of lines hitting a convex body K and having speed limits no slower than $v_0 > 0$:

$$\{(\ell, v) : \ell \in [K] \text{ and } v \geq v_0\}.$$

This has finite intensity measure, since $\gamma > 1$. It follows that the full set of marked lines $\{(\ell, v) : \ell \in [K]\}$ can be expressed as a countable union of random closed sets

(indeed, locally finite unions of lines) when broken up according to different ranges of speed limit. Hence, the union of all these lines does have a natural representation as a random F_σ. Indeed it can be related to random closed set theory as follows.

DEFINITION 1.3. For given $d \geq 2$ and $\gamma > 1$, and fixed $v_0 > 0$, let Π_{v_0} denote the *proper* marked Poisson line process of all lines with speed limits no slower than v_0:

$$\Pi_{v_0} = \{(\ell, v) \in \Pi : v \geq v_0\}.$$

The *silhouette* $\mathcal{S}_{v_0}$ of Π_{v_0} is the random closed set which is the union of all lines in Π_{v_0}:

$$\mathcal{S}_{v_0} = \bigcup\{\ell : (\ell, v) \in \Pi_{v_0}\} = \bigcup\{\ell : (\ell, v) \in \Pi \text{ and } v \geq v_0\}. \tag{1.9}$$

So $\mathcal{S}_{0+} = \bigcup_{v_0>0} \mathcal{S}_{v_0}$ can be viewed as a random F_σ.

Note that almost surely the *unmarked* line process $\{\ell : (\ell, v) \in \Pi_{v_0}\}$ can be recovered from the silhouette $\mathcal{S}_{v_0}$. Moreover, we can recover Π in entirety from the details of the changes in $\mathcal{S}_v$ as v varies, since $\mathcal{S}_v \setminus \mathcal{S}_{v+}$ is the locally finite union of lines whose speed limits are exactly equal to v. Indeed, $\Pi_v = \Pi_{v-} = \bigcap_{v<v_0} \Pi_v$ for $v_0 > 0$, and Π_v changes only at countably many values of $v > 0$, and, almost surely, for all $v > 0$ if $\mathcal{S}_v \setminus \mathcal{S}_{v+}$ is nonempty then it is composed of just one line. Thus,

$$\Pi = \{(v, \ell) : \mathcal{S}_v \setminus \mathcal{S}_{v_+} = \{\ell\} \text{ for some } v > 0\}.$$

For notational convenience, we introduce the *maximum speed limit function* V holding everywhere on $\mathbb{R}^d$ and imposed by Π. This is a random upper-semicontinuous function $V : \mathbb{R}^d \to [0, \infty)$ defined in terms of its upper level sets:

$$\{x : V(x) \geq v_0\} = \mathcal{S}_{v_0} = \bigcup\{\ell : (\ell, v) \in \Pi_{v_0}\}. \tag{1.10}$$

As with random dense line patterns, there is no satisfactory theory for general random upper-semicontinuous functions: we use V merely as a convenient mathematical shorthand for the filtration of random closed sets $\{\mathcal{S}_v : v > 0\}$.

2. Lipschitz paths and networks. This section introduces the notion of Π-paths; locally Lipschitz paths traversing a system of "roads" (supplied by Π) furnished with speed limits via the maximum speed limit function V. We will formulate this notion carefully and prove results yielding a variational context within which to study the minimum-time Π-paths ("Π = geodesics") between specified points. Care is needed, because we cannot assume that such paths are built up using consecutive sequences of intervals spent on different roads (and indeed this absolutely cannot be the case for dimension $d > 2$). Instead the Π-paths are best

viewed using such intervals arranged in a tree pattern rather than ordered sequentially (compare the use of trees to represent bounded variation paths in Hambly and Lyons [16]).

From henceforth, we shall fix a dimension $d \geq 2$ (since the case $d = 1$ is trivial) and a parameter $\gamma > 1$ (note however that the discussion of this section will lead to imposition of progressively more severe constraints on γ). Recall (e.g., from Evans [14], Chapter 5) that a Lipschitz curve $\xi : [0, T) \to \mathbb{R}^d$ satisfies $|\xi(s) - \xi(t)| \leq A|s - t|$ when $0 \leq s < t < T$, for some constant $A \geq 0$. The least such A is the *Lipschitz constant* $\mathrm{Lip}(\xi)$. The Lipschitz property for ξ using constant $\mathrm{Lip}(\xi) = A$ holds if and only if ξ is absolutely continuous with almost everywhere defined weak derivative ξ', with $\operatorname{ess\,sup}_t |\xi'(t)| \leq A$.

We first discuss two preparatory results; a simple lemma relating the velocity of a general Lipschitz path (defined for almost all time) to the directions of lines which it visits, and a corollary concerning the way in which such a Lipschitz path behaves at the intersections formed by a pattern of lines. Intuitively speaking, if the path velocity has nonzero component normal to a given line then it must move away immediately, so for almost all time either the path runs on the line and has velocity parallel to the line, or the path does not lie on the line at all.

LEMMA 2.1. *Let $\xi : [0, T) \to \mathbb{R}^d$ be a locally Lipschitz path and let ℓ be a line, both lying in d-dimensional space $\mathbb{R}^d$. Suppose that $\underline{e}$ is a unit vector parallel to the direction of ℓ. Then the time-set*

$$\{t \in [0, T) : \xi(t) \in \ell \text{ and } \xi'(t) \neq \langle \xi'(t), \underline{e} \rangle \underline{e}\}$$

is a Lebesgue-null subset of $[0, T)$.

PROOF. Let P denote projection onto the hyperplane normal to ℓ. The line ℓ projects under P to a singleton point set which we denote by $\{P\ell\}$. Restricting to compact subsets of $[0, T)$ if necessary, it suffices to treat the case in which ξ is globally Lipschitz over $[0, T)$; let $\alpha = \mathrm{Lip}(\xi)$ be the corresponding Lipschitz constant, so that

$$|\xi'(t)| \leq \alpha \qquad \text{for almost all } t \in [0, T).$$

The set $\{t \in [0, T) : P\xi(t) = P\ell \text{ and } P\xi'(t) \neq 0\}$ is the countable union of time-sets

$$A^{\pm}_{j,n} = \left\{ t \in [0, T) : P\xi(t) = P\ell \text{ and } \pm \langle P(\xi')(t), \underline{e}_j \rangle > \frac{1}{n} \right\} \qquad \text{for } j = 1, \ldots, d-1,$$

where $\underline{e}_1, \ldots, \underline{e}_{d-1}$ are orthogonal unit vectors perpendicular to ℓ, so it suffices to show each $A^{\pm}_{j,n}$ is Lebesgue-null (note that $P\xi'$ is not continuous, so that $A^{\pm}_{j,n}$ may not be open).

Without loss of generality, fix attention on $A^+_{1,n}$; we will complete the proof by showing that this is Lebesgue-null. Fix arbitrary $\varepsilon > 0$ and cover $A^+_{1,n}$ by a disjoint countable union of closed bounded intervals $F = \bigcup_i [a_i, b_i]$, such that

$$\text{Leb}(F \setminus A^+_{1,n}) < \frac{\varepsilon}{\alpha n}. \tag{2.1}$$

Since ξ is continuous, we may shrink each interval $[a_i, b_i]$ so as to ensure that $P\xi = P\ell$ at a_i and b_i, while still preserving (2.1) and the covering property $A^+_{1,n} \subseteq F$. Since $P\xi = P\ell$ on the end-points of each $[a_i, b_i]$,

$$\begin{aligned}\int_{F \setminus A^+_{1,n}} P\xi'(t)\,\mathrm{d}t + \int_{A^+_{1,n}} P\xi'(t)\,\mathrm{d}t &= \int_F P\xi'(t)\,\mathrm{d}t \\ &= \sum_i \bigl(P\xi(b_i) - P\xi(a_i)\bigr) = 0.\end{aligned}$$

However, we can apply the Lipschitz property of ξ to show that

$$\begin{aligned}\alpha\,\text{Leb}(F \setminus A^+_{1,n}) &\geq \left|\int_{F \setminus A^+_{1,n}} P\xi'(t)\,\mathrm{d}t\right| \geq \left|\int_{A^+_{1,n}} \langle P\xi'(t), \underline{e}_1\rangle\,\mathrm{d}t\right| \\ &\geq \frac{1}{n}\,\text{Leb}(A^+_{1,n})\end{aligned}$$

and, therefore,

$$\text{Leb}(A^+_{1,n}) \leq \alpha n\,\text{Leb}(F \setminus A^+_{1,n}) \leq \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, the result follows. □

As a direct consequence of Lemma 2.1, the only way in which a Lipschitz path can spend positive time on the intersection of two distinct lines is by being essentially at rest.

COROLLARY 2.2. *Consider a network in $\mathbb{R}^d$ formed by the union of a countable number of distinct lines $\ell_1, \ell_2, \ldots$, and form the* intersection point pattern

$$\mathcal{I} = \bigcup_{1 \leq i < j < \infty} \ell_i \cap \ell_j.$$

If $\xi : [0, T) \to \mathbb{R}^d$ is a locally Lipschitz curve in $\mathbb{R}^d$ then the time-set

$$\bigl\{t \in [0, T) : \xi(t) \in \mathcal{I} \text{ and } |\xi'(t)| > 0\bigr\} \tag{2.2}$$

must be a Lebesgue-null subset of $[0, T)$.

PROOF. Since $\mathcal{I}$ is a countable union of points, it suffices to consider two distinct lines ℓ_1 and ℓ_2 with nonempty intersection. Let $\underline{e}_i$ be a unit vector parallel to the direction of ℓ_i. Note that, since the ℓ_i are distinct and meet (and therefore

cannot be parallel), it follows that the unit vectors $\underline{e}_1$ and $\underline{e}_2$ must be linearly independent. By Lemma 2.1, the following two time-sets are both Lebesgue-null:

$$\{t \in [0,T) : \xi(t) \in \ell_1 \text{ and } \xi'(t) \neq \langle \xi'(t), \underline{e}_1 \rangle \underline{e}_1\},$$
$$\{t \in [0,T) : \xi(t) \in \ell_2 \text{ and } \xi'(t) \neq \langle \xi'(t), \underline{e}_2 \rangle \underline{e}_2\}.$$

If ξ' is simultaneously parallel to the two linearly independent $\underline{e}_i$ then it must vanish. Consequently, the time-set defined by equation (2.2) is a subset of the union of these two sets, and is therefore Lebesgue-null. □

We now define the fundamental notion explored in this paper.

DEFINITION 2.3. A Π-*path* is a locally Lipschitz path in $\mathbb{R}^d$ which for almost all time obeys the speed limits imposed by Π via the random upper semicontinuous function $V : \mathbb{R}_d \to [0, \infty)$. Expressed more precisely, the Π-path is given by an $\mathbb{R}^d$-valued function

$$\xi : [0,T) \to \mathbb{R}^d$$

defined up to a (possibly infinite) *terminal time* $T > 0$, and satisfying the condition that, for all $v > 0$, the time-set

$$\{t \in [0,T) : |\xi'(t)| > v \text{ and } \xi(t) \notin \mathcal{S}_v\}$$

is a Lebesgue-null subset of $[0, \infty)$. Let $\mathcal{A}_T$ be the space of all Π-paths defined up to time $T > 0$, and let $\mathcal{A} = \bigcup_{T>0} \mathcal{A}_T$ be the space of all Π-paths whatsoever.

REMARK 2.4. Direct arguments using Lemma 2.1 show that the Lebesgue-null condition in Definition 2.3 can be replaced by any one of three equivalent conditions:

1. for almost all $t \in [0,T)$ for which $\xi'(t) \neq 0$, it is the case that $\xi(t) \in \mathcal{S}_{|\xi'(t)|}$;
2. for almost all $t \in [0,T)$ for which $\xi'(t) \neq 0$, it is the case that the line $\xi(t) + \xi'(t)\mathbb{R}$ belongs to $\Pi_{|\xi'(t)|}$;
3. for almost all $t \in [0,T)$, it is the case that $|\xi'(t)| \leq V(\xi(t))$.

Here, $\xi(t) + \xi'(t)\mathbb{R}$ denotes the line through $\xi(t)$ with orientation $\xi'(t)$.

Condition 1 implies that for almost all $t \in [0,T)$, if $\xi(t) \notin \bigcup_{v>0} \mathcal{S}_v$ then $\xi'(t) = 0$.

Crucially, Π-paths integrate the measurable orientation field provided by Π and obey the speed limits imposed by Π.

LEMMA 2.5. *Suppose that ξ is a Π-path. For any $(\ell, v) \in \Pi$, let $\underline{e}$ be a unit vector parallel to the direction of ℓ. Then the following time-sets are both Lebesgue-null*:

$$\{t \in [0,T) : \xi(t) \in \ell, \xi'(t) \neq \langle \xi'(t), \underline{e} \rangle \underline{e}\}, \tag{2.3}$$

$$\{t \in [0,T) : \textit{for some } (\ell, v) \in \Pi, \xi(t) \in \ell, |\xi'(t)| > v\}. \tag{2.4}$$

PROOF. It is a direct consequence of Lemma 2.1 that the time-set at (2.3) is Lebesgue-null.

As for the time-set at (2.4), note that, for some Lebesgue null-set $\mathcal{N}_1$,

$$\begin{aligned}&\{t \in [0,T) : |\xi'(t)| > v \text{ and } \xi(t) \in \ell \text{ for some } (\ell, v) \in \Pi\} \\ &\quad \subseteq \{t \in [0,T) : \xi(t) \in \ell, \xi(t) \in \widetilde{\ell} \text{ for some } (\ell, v), (\widetilde{\ell}, w) \in \Pi \text{ with } w > v\} \\ &\quad\quad \cup \mathcal{N}_1\end{aligned}$$

(use the equivalent Π-path Definition 2.3). But this implies that

$$\begin{aligned}&\{t \in [0,T) : |\xi'(t)| > v \text{ and } \xi(t) \in \ell \text{ for some } (\ell, v) \in \Pi\} \\ &\quad \subseteq \{t \in [0,T) : |\xi'(t)| > v \text{ and } \xi(t) \in \mathcal{I}\},\end{aligned}$$

where $\mathcal{I}$ is the intersection point pattern derived from Π as in the statement of Corollary 2.2. By Corollary 2.2, the time-set on the right-hand side is Lebesgue-null. □

As noted at the start of this section, there are two related reasons for taking this rather abstract approach to Π-paths, as opposed to working only with paths built up sequentially from segments of lines in Π. First, in dimension $d > 2$ there are no nontrivial sequential Π-paths, since almost surely no two lines of Π will intersect. Second, even in the planar case of $d = 2$ we must consider nonsequential Π-paths as possible limits of simple Π-paths, for example, when establishing the existence of minimizers of Π-path functionals (specifically, the functional yielding accumulated elapsed time).

Here and in the following, we establish a number of statements about $\mathcal{A}_T$ and $\mathcal{A}$, all of which should be qualified as holding "almost-surely," since they depend on the random construction of the marked line process Π. Similarly, "constants" are actually random variables measurable with respect to the line process Π. Borrowing the terminology of random walks in random environments, it might be said that we are concerned with the quenched law based on the random environment provided by Π.

It is immediate from the definition that Π-paths are individually fully Lipschitz over time intervals in which the Π-path in question belongs to a specified compact set. As a consequence, we can establish a priori bounds on the length of any Π-path beginning in a specified compact set K and with terminal time at most $T < \infty$ (hence finite), so long as $\gamma \geq d$. These bounds will depend on $T < \infty$, K, γ, and also on the random marked line pattern Π, and will allow us to deduce a uniform Lipschitz property for all Π-paths from $\mathcal{A}_T$ which start in K. Moreover, if $\gamma \geq d$ then the set $\mathcal{A}_T$ is weakly closed and, therefore, the family of all Π-paths from $\mathcal{A}_T$ started in K is weakly compact. This allows us to make sense of the notion of Π-*geodesics*, measuring path length not as geometric length but using total time of passage along the path. We first establish an a priori upper bound for the Lipschitz

constants of Π-paths begun near the origin and defined up to a finite time. (A more quantitative bound could be obtained by applying the ideas of the more refined comparison analysis of Theorem 5.2.)

THEOREM 2.6 (An a priori bound for path space). *Suppose that $\gamma \geq d \geq 2$. Fix $T < \infty$ and $r_0 > 0$, and consider a Π-path ξ with $\xi(0) \in \mathrm{ball}(\mathbf{o}, r_0)$. Then there is $C = C(\gamma, T, r_0, \Pi) < \infty$ (which is a constant when conditioned on the particular realization of the marked line process Π) such that almost surely*

$$\big|\xi(t)\big| \leq C \qquad \textit{for all } t \in [0, T).$$

PROOF. We consider a given realization of the speed limit function V (equivalently, of Π). It suffices to obtain a lower bound on the time at which ξ first exits $\mathrm{ball}(\mathbf{o}, r)$ for a given $r > r_0$, and to show that this bound tends to infinity as $r \to \infty$. We achieve this by constructing a comparison with a one-dimensional system (controlled by the speed limit function V), thus delivering an upper bound on $|\xi|$, and then by showing that with probability 1 the underlying configuration of Π is such that the comparison system takes infinite time to diverge to infinity.

Our comparison system $y : [0, \infty) \to [0, \infty)$ satisfies

$$\begin{aligned} y'(t) &= \overline{V}\big(y(t)\big) = \sup\big\{V(x) : |x| \leq y(t)\big\} \\ &= \sup\big\{v : (\ell, v) \in \Pi, \ell \Uparrow \mathrm{ball}(\mathbf{o}, r)\big\}, \\ y(0) &= r_0. \end{aligned} \tag{2.5}$$

So y could be thought of as the distance from $\mathbf{o}$ of an idealized path which always travels radially at the maximum speed $\overline{V}$ available from Π within that distance. Standard comparison techniques show that if ξ is a Π-path then

$$y(t) \geq \big|\xi(t)\big| \qquad \text{for almost all } t \text{ whenever } \big|\xi(0)\big| \leq r_0.$$

The next step is to control the growth of the maximum speed limit $\overline{V}(y)$ as a function of y. We introduce a nonlinear projection from marked line-space to the quadrant $[0, \infty)^2$:

$$(\ell, v) \mapsto \big(|r|^{d-1}, v^{-(\gamma-1)}\big) = (p, s).$$

We think of $p = |r|^{d-1}$ as "generalized distance," and $s = v^{-(\gamma-1)}$ as "meta-slowness." The fibers of the projection have zero invariant line measure. Bearing in mind that $\gamma \geq d \geq 2$, and using the expression for the intensity measure at (1.8), also for μ_d at (1.5), the image of Π under the projection is a stationary Poisson point process on the quadrant $[0, \infty)^2$, with intensity measure

$$\begin{aligned} &(\gamma - 1)\frac{\omega_{d-1}}{2} \times \big(-v^{-\gamma}\,\mathrm{d}v\big) \times \mathrm{d}\big(|r|^{d-1}\big) \\ &\quad = \frac{\omega_{d-1}}{2}\,\mathrm{d}\big(v^{-(\gamma-1)}\big)\,\mathrm{d}\big(|r|^{d-1}\big) = \frac{\omega_{d-1}}{2}\,\mathrm{d}p\,\mathrm{d}s \end{aligned}$$

(recall that ω_{d-1} is the hyper-surface area of the unit ball in $\mathbb{R}^d$).

This projection to a planar point process delivers a Poisson point process $\widetilde{\Pi}$ of constant intensity $\frac{1}{2}\omega_{d-1}$ in the quadrant. Note that, for $r > 0$,

$$\big(\overline{V}(r)\big)^{-(\gamma-1)} = \inf\big\{s : (p, s) \in \widetilde{\Pi} \text{ and } p \le r^{d-1}\big\}.$$

Accordingly, we can use (2.5) to obtain information about the sequence of generalized distances at which the meta-slowness of the comparison system y changes, deriving a recursive expression (compare the use of perpetuities in [17, 19]). Initially, the random meta-slowness of y at generalized distance $P_0 = r_0^{d-1}$ is

$$S_0 \text{ of distribution Exponential}\big(\tfrac{1}{2}\omega_{d-1}P_0\big).$$

Let $P_0 < P_1 < P_2 < \cdots$ be the sequence of generalized distances at which the meta-slowness changes through values $S_0 > S_1 > S_2 > \cdots$. Poisson process arguments show

$$\begin{aligned} P_n - P_{n-1} &= \frac{1}{S_{n-1}}T_n \qquad \text{for } T_n \text{ of distribution Exponential}\Big(\frac{1}{2}\omega_{d-1}\Big); \\ S_n &= U_n S_{n-1} \qquad \text{for } U_n \text{ of distribution Uniform}(0,1). \end{aligned} \tag{2.6}$$

Moreover, the T_i's and U_i's are all independent.

It is useful to reorganize this recursion so as to recognize the product $X_n = S_n P_n$ as a Markov chain (in fact, a perpetuity):

$$X_n = S_n P_n = U_n(T_n + S_{n-1}P_{n-1}) = U_n(T_n + X_{n-1}). \tag{2.7}$$

Iteration shows that this perpetuity has a limiting distribution, expressible as an infinite almost-surely convergent sum

$$U_1T_1 + U_1U_2T_2 + U_1U_2U_3T_3 + \cdots$$

(stronger results on perpetuities can be found in the foundational paper of Vervaat [25]). Indeed, Markov chain arguments show that X_n converges to this distribution in total variation with geometric ergodicity [22], Chapter 15, especially Theorem 15.0.1. This follows by noting that the Markov chain $\{X_n : n \ge 0\}$ is Lebesgue-irreducible and satisfies a geometric Foster–Lyapunov condition based on (e.g.)

$$\begin{aligned} &\mathbb{E}[1 + X_n | X_{n-1}] \\ &\quad = 1 + \frac{1}{2}\Big(\frac{2}{\omega_{d-1}} + X_{n-1}\Big) \\ &\quad \le \begin{cases} \dfrac{2}{3}(1 + X_{n-1}), & \text{if } 1 + X_{n-1} > 3\Big(1 + \dfrac{2}{\omega_{d-1}}\Big), \\ \dfrac{2}{3}(1 + X_{n-1}) + \dfrac{1}{2} + \dfrac{1}{\omega_{d-1}}, & \text{otherwise.} \end{cases} \end{aligned}$$

The interval $[0, 3(1 + \frac{2}{\omega_{d-1}}))$ is a small set for the Markov chain X (since U_n and T_n are independent and have continuous densities which are strictly positive over their ranges of $[0, 1]$ and $[0, \infty)$), so this is indeed a geometric Foster–Lyapunov condition, and establishes geometric ergodicity for the Markov chain X.

Consider the elapsed actual time between generalized positions P_{n-1} and P_n, namely

$$(S_{n-1})^{1/(\gamma-1)} \times \left(P_n^{1/(d-1)} - P_{n-1}^{1/(d-1)}\right).$$

Summing over n, and using the requirement that $\gamma \geq d$, the total time until y reaches infinity is given by the sum

$$\begin{aligned}
&\sum_n (S_{n-1})^{1/(\gamma-1)} \times \left(P_n^{1/(d-1)} - P_{n-1}^{1/(d-1)}\right) \\
&\quad = \sum_n (S_{n-1})^{1/(\gamma-1)-1/(d-1)} \times \left((S_{n-1}P_n)^{1/(d-1)} - (S_{n-1}P_{n-1})^{1/(d-1)}\right) \\
&\quad = \sum_n (S_{n-1})^{1/(\gamma-)1-1/(d-1)} \\
&\qquad\qquad \times \left((S_{n-1}P_{n-1} + T_n)^{1/(d-1)} - (S_{n-1}P_{n-1})^{1/(d-1)}\right).
\end{aligned}$$

If $\gamma = d$, then $(S_{n-1})^{1/(\gamma-1)-1/(d-1)} = 1$. If on the other hand $\gamma > d$, then the exponent $\frac{1}{\gamma-1} - \frac{1}{d-1}$ is negative, while from the recurrence it follows that $S_n \to 0$ almost surely. In either case, the above sum diverges if the same can be said for

$$\sum_n \left((S_{n-1}P_{n-1} + T_n)^{1/(d-1)} - (S_{n-1}P_{n-1})^{1/(d-1)}\right). \tag{2.8}$$

However, geometric ergodicity of $X_n = S_n P_n$, together with independence of the T_i's and U_i's, and the exponential distribution of T_n, allows us to deduce that almost surely infinitely many of these summands must satisfy

$$(S_{n-1}P_{n-1} + T_n)^{1/(d-1)} - (S_{n-1}P_{n-1})^{1/(d-1)} > 1.$$

Hence, it follows that the infinite sum (2.8) almost surely diverges and, therefore, y will almost surely takes infinite time to reach infinity. This suffices to prove the theorem. □

REMARK 2.7. It can be shown that in case $\gamma < d$ then there are indeed Π-paths which reach infinity in finite time, and even in finite expected time.

REMARK 2.8. We will see (Theorem 3.1) that if we require $\gamma > d$ then in addition we can connect specified pairs of points by Π-paths in finite time.

We can now deduce the existence of an a priori bound on the global Lipschitz constant of Π-paths of finite length begun in a fixed compact set. For any constant C, the lines from Π meeting ball($\mathbf{o}$, C) have speed-limits bounded above by

$\overline{V}(C)$, a random value depending on the random environment Π. Fixing $T < \infty$ and $r_0 > 0$, and taking C to be the random value depending on Π in the statement of Theorem 2.6, then Remark 2.4 condition 3 implies that the Π-paths in $\mathcal{A}_T$ beginning in ball$(\mathbf{o}, r_0)$ satisfy a uniform Lipschitz property with random Lipschitz constant depending on Π, r_0, and T. Note that $r_0 > 0$ can be chosen arbitrarily.

A closure result for $\mathcal{A}_T$ follows immediately. Recall (again using results expounded in [14], Chapter 5) that the Sobolev space $W^{1,2}([0,T) \to \mathbb{R}^d)$ can be viewed as the space of absolutely continuous curves $\xi : [0,T) \to \mathbb{R}^d$ whose first derivatives ξ' satisfy $\int_0^T |\xi'|^2 \, \mathrm{d}t < \infty$. Thus $W^{1,2}([0,T) \to \mathbb{R}^d)$ forms a separable Hilbert space when furnished with an inner-product norm given by $\sqrt{(\int_0^T |\xi|^2 \, \mathrm{d}t + \int_0^T |\xi'|^2 \, \mathrm{d}t)}$. Recall also the Hilbert-space fact that bounded and weakly-closed subsets of $W^{1,2}([0,T) \to \mathbb{R}^d)$ are weakly compact.

LEMMA 2.9 (Closure of path space). *Suppose $\gamma \geq d$. For finite T, the path space $\mathcal{A}_T$ is closed in the Sobolev space $W^{1,2}([0,T) \to \mathbb{R}^d)$.*

PROOF. Consider any sequence $\xi_1, \xi_2, \ldots$ of paths drawn from $\mathcal{A}_T$. Suppose $\xi_n \to \xi$ when considered as elements of $W^{1,2}([0,T) \to \mathbb{R}^d)$. By Sobolev space arguments, taking a suitably convergent subsequence if necessary, we may suppose that

$$\begin{aligned} \xi_n(t) &\to \xi(t) \qquad \text{for all } t \in [0,T); \\ \xi_n'(t) &\to \xi'(t) \qquad \text{for almost all } t \in [0,T). \end{aligned}$$

Consider the following time-set, which can be seen to be Lebesgue-null by combining the Lipschitz property of ξ, the choice of the convergent sequence, and using the equivalent definition of Π-paths given by Remark 2.4 condition 2:

$$\begin{aligned} \mathcal{N} = \{ t : \text{ either } \xi'(t) \text{ does not exist, or } \xi_n'(t) \not\to \xi'(t), \\ \text{or, for some } n, \xi_n'(t) \neq 0 \text{ but } \xi_n(t) + \xi_n'(t)\mathbb{R} \notin \mathcal{S}_{|\xi_n'(t)|} \}. \end{aligned}$$

Fix attention on $t \notin \mathcal{N}$ for which $|\xi'(t)| > u > 0$. Then $\xi_n'(t) \to \xi'(t)$, so

$$\xi_n(t) + \xi_n'(t)\mathbb{R} \in \Pi_u \qquad \text{for all large enough } n.$$

Furthermore, since Π_u yields a locally finite line process,

$$\begin{aligned} &\text{either } \xi_n(t) + \xi_n'(t)\mathbb{R} \text{ eventually equals } \ell_t \text{ for some fixed } \ell_t \text{ from } \Pi_u \\ &\text{and moreover } \xi(t) + \xi'(t)\mathbb{R} = \ell_t, \\ &\text{or } \xi_n(t) \to \text{ an intersection point of the line process } \Pi_u. \end{aligned}$$

In the first case, since ℓ_t is a line from $\Pi_{|\xi_n'(t)|}$ for all large enough n, it follows from the convergence $\xi_n'(t) \to \xi'(t)$ that ℓ_t is a line from $\Pi_{|\xi'(t)|}$.

In the second case, Corollary 2.2 implies that

$$\{ t : \xi(t) \text{ is an intersection point of } \Pi \text{ and } \xi'(t) \neq 0 \}$$

must be a null time-set. Either way, up to a Lebesgue-null time-set,

either $\xi'(t) = 0$,
or $\xi(t) + \xi'(t)\mathbb{R} \in \Pi_{|\xi'(t)|}$.

Hence, ξ is a Π-path (using Remark 2.4, specifically the equivalent defining condition 2 for Π-paths), hence belongs to $\mathcal{A}_T$. $\square$

We can improve on this lemma to show that $\mathcal{A}_t$ is *weakly closed*, from which there follows a useful compactness result.

THEOREM 2.10 (Weak closure of path space). *Suppose $\gamma \geq d$. For finite T, the path space $\mathcal{A}_T$ is weakly closed in $W^{1,2}([0,T) \to \mathbb{R}^d)$.*

PROOF. Consider any sequence $\xi_1, \xi_2, \ldots$ in $\mathcal{A}_T$. Suppose $\xi_n \to \xi$ *weakly* in $W^{1,2}([0,T) \to \mathbb{R}^d)$. We will invoke the equivalent defining condition 3 from Remark 2.4, namely, if we can show that $|\xi'(t)| \leq V(\xi(t))$ for almost all $t \in [0,T)$ then ξ is a Π-path. As an immediate consequence, this will imply that $\mathcal{A}_T$ is weakly closed.

If $T < \infty$, then weak convergence in $W^{1,2}([0,T) \to \mathbb{R}^d)$ implies uniform convergence; in particular this implies ξ is continuous.

Fix $v > 0$, and consider a nonempty time interval $[r,s] \subseteq [0,T)$, such that $[r,s]$ is contained in a single connected component of the open time-set $\{t \in [0,T) : \xi(t) \notin \mathcal{S}_v\}$. Thus the compact set $\{\xi(t) : r \leq t \leq s\}$ does not intersect the closed set $\mathcal{S}_v$ and indeed will not intersect the closed set $\mathcal{S}_{v-\varepsilon}$ for some (perhaps small) $\varepsilon \in (0,v)$. By the uniform convergence of ξ_n to ξ, we know that for all large enough n the compact set $\{\xi_n(t) : r \leq t \leq s\}$ does not intersect $\mathcal{S}_{v-\varepsilon}$; since ξ_n is a Π-path, this implies that $|\xi_n'(t)| < v - \varepsilon$ for almost all $t \in [r,s]$, and hence that ξ_n is Lipschitz with Lipschitz constant $v - \varepsilon$ over the time interval $[r,s]$. But the uniform convergence of ξ_n to ξ implies that ξ itself is Lipschitz with Lipschitz constant $v - \varepsilon$ over the time interval $[r,s]$, therefore, that $|\xi'(t)| \leq v - \varepsilon < v$ for almost all $t \in [r,s]$.

Expressing the open time-set $\{t \in [0,T) : \xi(t) \notin \mathcal{S}_{v-\varepsilon}\}$ as a countable union of closed intervals, and letting $\varepsilon \to 0$, it follows that for almost all $t \in [0,T)$ if $\xi(t) \notin \mathcal{S}_v$ then $|\xi'(t)| < v$. Consequently, $|\xi'(t)| \leq V(\xi(t))$ for almost all $t \in [0,T)$ as required. $\square$

It follows immediately that bounded subsets of $\mathcal{A}_T$ are weakly precompact, and this is crucial for the discussion of minimum-time Π-paths given in the next section.

COROLLARY 2.11 (Weak compactness and path space). *Suppose $\gamma \geq d$. Fix $T > 0$ and consider the set of paths in $\mathcal{A}_T$ which begin in a compact set K. This set is weakly compact as a subset of $W^{1,2}([0,T) \to \mathbb{R}^d)$.*

PROOF. Immediate from Theorems 2.6 and 2.10 together with weak compactness results for the Hilbert space $W^{1,2}([0,T) \to \mathbb{R}^d)$. $\square$

3. Π-paths between arbitrary points. We have shown that the condition $\gamma \geq d$ ensures that almost surely Π-paths do not diverge to infinity in finite time (as noted in Remark 2.7, this condition is also necessary). However, it is not yet clear whether Π-paths can move from one point to another in finite time. In dimension $d = 2$, the question is essentially one of whether one can reach a specified point in finite time by travelling along paths of progressively slower and slower speed. Dimension $d \geq 3$ appears intransigent at first glance, since one knows that lines of a Poisson line process do not intersect each other in dimension 3 and higher. Nevertheless, almost surely there are Π-paths connecting *all* pairs of points in dimension 2 *and* higher, so long as we strengthen the condition on γ to $\gamma > d$.

We begin by showing how to connect specified pairs of points.

THEOREM 3.1 (Π-paths connect given pairs of points in finite time). *Suppose $\gamma > d$. Then specified x_1 and x_2 in $\mathbb{R}^d$ can almost surely be connected in finite time T by a Π-path ξ.*

PROOF. The construction connects segments of lines from Π together in a tree-like fashion, rather than sequentially. The basic idea is as follows: for a sufficiently large constant α (indeed, $\alpha > 2^{(\gamma-1)/(\gamma-d)}$ suffices), construct disjoint balls of radius $|x_1 - x_2|/\alpha$ around x_1 and x_2. Choose the fastest line ℓ in Π hitting both balls, corresponding to the root of a binary tree representation of a path connecting x_1 to x_2. Then create two daughter nodes, repeating the construction based on (a) x_1 and the closest point to x_1 on ℓ, and (b) x_2 and the closest point to x_2 on ℓ. Extend this recursively to generate a binary tree-indexed collection of line segments. Figure 2 illustrates the first two stages of the $d = 2$ case.

The path ξ formed from this binary tree is evidently a Π-path. The issue is to show that it makes the connection from x_1 to x_2 in finite time.

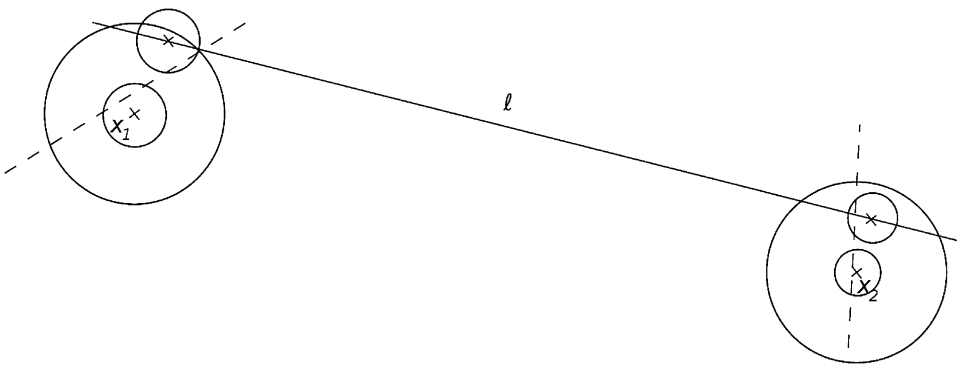

FIG. 2. *First two stages of a recursive construction of a Π-path from x_1 to x_2 in two dimensions, using fastest available lines taken from an improper Poisson line process marked by speeds. Note that in the case of higher dimensions the lines will almost surely not intersect.*

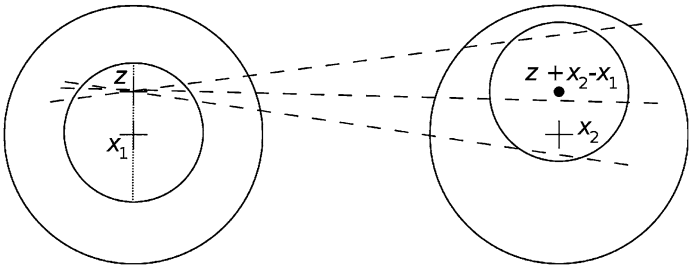

FIG. 3. *Reduction of* $[\text{ball}(x_1, \alpha^{-1}r)] \cap [\text{ball}(x_2, \alpha^{-1}r)]$ *to a smaller hitting set for which the line measure is more easily computable yet still provides a useful lower bound.*

First, we need a stochastic lower bound for the speed limit of the fastest line connecting two balls, namely the speed of the fastest line of Π in

$$\big[\text{ball}(x_1, \alpha^{-1}r)\big] \cap \big[\text{ball}(x_2, \alpha^{-1}r)\big].$$

Here, we write $r = |x_1 - x_2|$ for the Euclidean distance between x_1 and x_2. We obtain a stochastic lower bound for the speed distribution in two steps:

(a) Shrink $\text{ball}(x_1, \alpha^{-1}r)$ to $\text{ball}(x_1, \alpha^{-1}r/2)$ and then replace the smaller $\text{ball}(x_1, \alpha^{-1}r/2)$ by the hyperdisk $D(x_1, \alpha^{-1}r/2)$ obtained by intersecting $\text{ball}(x_1, \alpha^{-1}r/2)$ with the hyperplane through x_1 which is normal to the vector $x_2 - x_1$;

(b) Consider the bundle of lines in $[D(x_1, \alpha^{-1}r/2)] \cap [\text{ball}(x_2, \alpha^{-1}r)]$ which run through a given point $z \in D(x_1, \alpha^{-1}r/2)$. For each such z, reduce the bundle size by restricting attention to lines which additionally intersect $\text{ball}(z + x_2 - x_1, \alpha^{-1}r/2) \subset \text{ball}(x_2, \alpha^{-1}r)$.

This geometric construction is illustrated in Figure 3.

Using inclusion of hitting sets, this produces an easily computable lower bound for the line measure:

$$\begin{aligned}
&\mu_d\big(\big[\text{ball}(x_1, \alpha^{-1}r)\big] \cap \big[\text{ball}(x_2, \alpha^{-1}r)\big]\big) \\
&\quad \geq \mu_d\big(\big[D(x_1, \alpha^{-1}r/2)\big] \cap \big[\text{ball}(x_2, \alpha^{-1}r)\big]\big) \\
&\quad \geq m_{d-1}\big(D(x_1, \alpha^{-1}r/2)\big) \times \mu_{d-1}^{(\mathbf{o})}\big(\big[\text{ball}(x_2 - x_1, \alpha^{-1}r/2)\big]\big).
\end{aligned}$$

Here, $\mu_{d-1}^{(\mathbf{o})} = \frac{\sin\theta}{\kappa_{d-1}} m_{S_+^{d-1}}$ is derived from the disintegration of μ_d by Lebesgue measure on the hyperplane through x_1 which is normal to $x_2 - x_1$, using (1.7). Thus $\mu_{d-1}^{(\mathbf{o})}$ is a weighted version of the invariant (hyper-surface area) measure on the hemisphere of unsensed lines ℓ passing through the origin $\mathbf{o}$, weighted by $\sin\theta$ where θ is the angle between ℓ and the hyperplane, normalized to have unit total measure.

Recall that κ_{d-1} denotes the $(d-1)$-volume of the unit $(d-2)$-ball. Thus,

$$m_{d-1}\big(D(\xi_1, \alpha^{-1}r/2)\big) = \kappa_{d-1}\left(\frac{r}{2\alpha}\right)^{d-1}.$$

On the other hand, from (1.7) the relevant computation of weighted hyper-surface area for the visibility hemisphere is

$$\frac{\omega_{d-2}}{\kappa_{d-1}} \int_{\pi/2-\theta_0}^{\pi_2} \sin\theta \cos^{d-2}\theta \, \mathrm{d}\theta = \frac{\omega_{d-2}}{(d-1)\kappa_{d-1}} \cos^{d-1}\theta_0 = \cos^{d-1}\theta_0,$$

where $\cos\theta_0 = 1/(2\alpha)$ [and noting that $\omega_{d-2} = (d-1)\kappa_{d-1}$]; hence

$$\mu_{d-1}^{(\mathbf{o})}\big([\mathrm{ball}(x_2 - x_1, \alpha^{-1}r/2)]\big) = \left(\frac{1}{2\alpha}\right)^{d-1}.$$

These considerations yield the lower bound

$$\mu_d\big([\mathrm{ball}(\xi_1, \alpha^{-1}r)] \cap [\mathrm{ball}(\xi_2, \alpha^{-1}r)]\big) \geq \kappa_{d-1}\left(\frac{r}{4\alpha^2}\right)^{d-1}.$$

This reasoning can be applied to the recursive construction indicated above. Let r_h be the distance between points at node h on the binary tree representing the path, then (omitting some implicit conditioning on r_h)

$$\mathbb{P}[\text{fastest line speed at } h \leq v_h] \leq \exp\left(-\frac{\kappa_{d-1}}{(4\alpha^2)^{d-1}} \frac{r_h^{d-1}}{v_h^{\gamma-1}}\right).$$

By construction, if node h is at level n of the tree then $r_n \leq \alpha^{-n} r_0$, where r_0 is the Euclidean distance between the original start and finish points. Fixing $\varepsilon > 0$, we set

$$v_h = \frac{r_h^{(d-1)/(\gamma-1)}}{(n\zeta)^{1/(\gamma-1)}} \qquad \text{where } \zeta = \frac{(4\alpha^2)^{d-1}}{\kappa_{d-1}}(1+\varepsilon)\log 2.$$

Then $r_h^{d-1}/v_h^{\gamma-1} = n\zeta$. Use the first Borel–Cantelli lemma, and the convergence of

$$\begin{aligned}\sum_h \exp\left(-\frac{\kappa_{d-1}}{(d-1)(4\alpha^2)^{d-1}} \frac{r_h^{d-1}}{v_h^{\gamma-1}}\right) &= \sum_n 2^n \exp\big(-(1+\varepsilon)n\log 2\big) \\ &= \sum_n 2^{-\varepsilon n} < \infty,\end{aligned}$$

to deduce that it is almost surely the case that the speed limits of all but finitely many segments h in the binary tree representation will exceed

$$v_h = \frac{r_h^{(d-1)/(\gamma-1)}}{(n\zeta)^{1/(\gamma-1)}}.$$

By the triangle inequality, the relevant path length for node h is no greater than $(1 + \frac{2}{\alpha})r_h$. So the total time spent traversing the path is finite when

$$\sum_h \left(1 + \frac{2}{\alpha}\right)\frac{r_h}{v_h} = \sum_n 2^n \times \left(1 + \frac{2}{\alpha}\right)(n\zeta)^{1/(\gamma-1)} \left(\frac{r_0}{\alpha^n}\right)^{(\gamma-d)/(\gamma-1)} < \infty.$$

But this sum converges when $\alpha > 2^{(\gamma-1)/(\gamma-d)}$: thus, in this case the construction gives a finite-time Π-path between x_1 and x_2. $\square$

REMARK 3.2. It can be shown that $\gamma > d$ is also a necessary condition for connection of x_1 to x_2 by a Π-path in $\mathbb{R}^d$. For if $\gamma = d$ then all Π-paths leaving the origin are subject to an upper bound using the comparison process of Theorem 2.6, and a direct calculation shows that this comparison process takes infinite expected time to leave the origin.

Note that, for $d > 2$, the finite-time path has a curious fractal-like property: whenever the Π-path changes from one line of positive speed limit to another, then it must shift gears right down to zero speed then right up again to the new speed. (And the same applies to each change of speed limit while shifting gears down, and so on *ad infinitum*, as in the case of the fleas and poets described by Swift [24], On Poetry: a Rhapsody.)

Since there are only countably many lines in Π, a simple modification of the above construction shows that almost surely all lines in Π are connected.

COROLLARY 3.3. *If $\gamma > d$, then almost surely all lines of Π are joined by finite-time Π-paths.*

If two points x_1 and x_2 are joined by Π-paths taking finite time, then it is reasonable to ask whether there is a minimum-time Π-path. As a consequence of Corollary 2.11, we know that this occurs, since $\gamma > d$, and hence *a fortiori* the compactness condition $\gamma \geq d$ holds. We summarize this conclusion by means of a definition and a further corollary. We note in passing that Π-geodesics inherit all the properties of minimal geodesics in metric spaces; for example a minimal Π-geodesic cannot intersect itself.

DEFINITION 3.4 (Π-geodesic). The Π-path $\xi : [0, T] \to \mathbb{R}^d$ is said to be a Π-*geodesic* (or *minimum-time geodesic*) from $\xi(0) = x_1$ to $\xi(T) = x_2$ if there are no Π-paths connecting x_1 and x_2 in $\mathcal{A}_S$ for $S < T$.

COROLLARY 3.5 (Existence of Π-geodesics). *Suppose $\gamma > d$. Consider the set of paths ξ in $\mathcal{A}$ which begin at fixed location $\xi(0) = x_1$ and end at fixed location $\xi(T) = x_2$ (here T depends on ξ). Almost surely there exist Π-geodesics in $\mathcal{A}$ from x_1 to x_2.*

PROOF. By Theorem 3.1, almost surely there are connecting Π-paths in $\mathcal{A}_T$ for large enough $T < \infty$. Consider a sequence of such paths $\xi_1, \xi_2, \ldots$, starting at x_1 and ending at x_2, such that $\xi_n(t) = x_2$ for all $t \in [T_n, T)$, and such that T_n tends to T_∞ the infimum of all connection times for Π-paths between x_1 and x_2. By Corollary 2.11, we may extract a weakly convergent subsequence, and the limit

$\xi \in \mathcal{A}_T$ will satisfy $\xi(t) = x_2$ for all $t \in [T_\infty, T)$ and hence realize the infimum. The resulting Π-path ξ will be a Π-geodesic between x_1 and x_2. $\square$

In the next section, we will examine the extent to which Π-geodesics are uniquely determined by their end points. Before turning to this matter, we improve on Theorem 3.1 by showing that if $\gamma > d$ then almost surely *all* pairs of points in $\mathbb{R}^d$ are connected by finite-time Π-paths. That is to say, almost surely there are no infinite singularities in the metric space induced by the time taken by fastest Π-path transit. The proof closely follows that of Theorem 3.1, but splits paths apart in a hierarchical way so as to access entire regions rather than single points.

THEOREM 3.6. *Suppose $\gamma > d$ and $d \geq 2$. With probability* 1, *the network formed by Π connects up all pairs of points in $\mathbb{R}^d$ using finite-time Π-paths.*

PROOF. It suffices to establish that, almost surely, finite-time Π-paths can be used to connect a specified point to all the points of a single hypercube of positive area.

Consider then the construction of a path ξ from x_1 to a hypercube centered on x_2, where $|x_1 - x_2| = r_0$ and the hypercube is of side-length $\alpha^{-1} r_0$ for some sufficiently large integer α [indeed, $\alpha > (1 + \frac{\sqrt{d}}{2}) 2^{(\gamma-1)/(\gamma-d)}$ suffices]. The construction commences as in Theorem 3.1, choosing the fastest line ℓ of Π in $[\text{ball}(x_1, \alpha^{-1} r_0)] \cap [\text{ball}(x_2, \alpha^{-1} r_0)]$, and this corresponds to the root of a tree now representing a whole family of paths. Repeat the construction, adding a further line from Π which nearly connects x_1 to the point on ℓ closest to x_1, as in Theorem 3.1. However, on the other side, we generate α^d separate paths, using the fastest possible line to connect the ball of radius $\alpha^{-1} r_0$ centered on the point on ℓ closest to x_2, to each of a total of α^d balls of radius $\alpha^{-1} r_0$ centered on centroids of cells arising from a dissection of the original hypercube of side-length $\alpha^{-1} r_0$ into α^d sub-hypercubes each of side-length $\alpha^{-2} r_0$. This is illustrated in Figure 4.

In the first case, the new distance is at most $\alpha^{-1} r_0$. In the second case, we may use Pythagoras to show that the new distance is at most $(1 + \frac{1}{2}\sqrt{d}) \alpha^{-1} r_0$. This

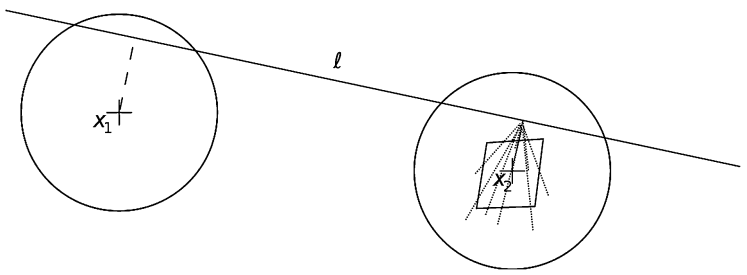

FIG. 4. *Initial stage of connecting x_1 to points in a specified hypercube* (*case of $d = 2$*).

construction step generates $1+\alpha^d$ new segments at the second level of the tree, the first one being like the segments generated in Theorem 3.1, while the remaining α^d nearly connect a given point to α^d centroids of sub-hypercubes.

Repeating the construction down to level n, we generate h_n segments of the first kind and k_n segments of the second kind, where

$$\begin{aligned} h_n &= 2h_{n-1}+k_{n-1}, \qquad h_1=1,\\ k_n &= \alpha^d k_{n-1}, \qquad k_1=\alpha^d. \end{aligned} \tag{3.1}$$

The total number of segments which have been built at level n is therefore

$$h_n+k_n = 2(h_{n-1}+k_{n-1})+\bigl(\alpha^d-1\bigr)k_{n-1}.$$

The bound on α imposed at the start of the proof, together with $\gamma > d \geq 2$, shows that $\alpha^d > 8$, and so we can use the recursion (3.1) and the fact that $h_1+k_1 = 1+\alpha^d$ to deduce

$$\begin{aligned} h_n+k_n &\leq 2(h_{n-1}+k_{n-1})+\bigl(\alpha^d-1\bigr)k_{n-1}\\ &= 2^2(h_{n-2}+k_{n-2})+\bigl(\alpha^d-1\bigr)(k_{n-1}+2k_{n-2})\\ &= \cdots = 2^{n-1}(h_1+k_1)+\bigl(\alpha^d-1\bigr)\bigl(k_{n-1}+2k_{n-2}+\cdots+2^{n-2}k_1\bigr)\\ &\leq \text{constant}\times\alpha^{nd}. \end{aligned} \tag{3.2}$$

Now consider the speed limit of the line forming node h at level n. Suppose the relevant distance between target points is r_h. Then (as in Theorem 3.1)

$$\mathbb{P}[\text{fastest line speed limit at } h \leq v_h] \leq \exp\biggl(-\frac{\kappa_{d-1}}{(4\alpha^2)^{d-1}}\frac{r_h^{d-1}}{v_h^{\gamma-1}}\biggr).$$

We know

$$r_h \leq \biggl(\frac{1+\sqrt{d}/2}{\alpha}\biggr)^n r_0.$$

Choose

$$v_h = \frac{r_h^{(d-1)/(\gamma-1)}}{(n\zeta)^{1/(\gamma-1)}} \qquad \text{as before, but with } \zeta = \frac{(4\alpha^2)^{d-1}}{\kappa_{d-1}}(1+\varepsilon)d\log\alpha.$$

Using the first Borel–Cantelli lemma once again, all but finitely many of the segments in this construction have speed limit exceeding the respective v_h, since

$$\sum_n \alpha^{dn}\exp\bigl(-(1+\varepsilon)nd\log\alpha\bigr) = \sum_n \alpha^{-\varepsilon nd} < \infty.$$

Thus each one of the paths can be traversed in finite time if the following sum converges:

$$\sum_{h \text{ in specified path}} \frac{r_h}{v_h} \leq \sum_n 2^n (n\zeta)^{1/(\gamma-1)} \left(\frac{1+\sqrt{d}/2}{\alpha} \right)^{n(\gamma-d)/(\gamma-1)}$$

$$\leq \zeta^{1/(\gamma-1)} \sum_n n^{1/(\gamma-1)} \left(2 \left(\frac{1+\sqrt{d}/2}{\alpha} \right)^{(\gamma-d)/(\gamma-1)} \right)^n.$$

Recalling the stipulation that $\gamma > d$, this converges if we choose

$$\alpha > \left(1 + \frac{\sqrt{d}}{2}\right) 2^{(\gamma-1)(\gamma-d)}.$$

The family of Π-paths used here is weakly compact (Corollary 2.11). It follows therefore that this construction almost surely delivers Π-paths which within finite time connect a specified point x_1 to all points in a nonempty hypercube with centroid x_2 and side length $\alpha^{-1}|x_2 - x_1|$. Using this fact together with judicious concatenation of Π-paths, it follows that almost surely all pairs of points in $\mathbb{R}^d$ are connected by Π-geodesics. □

In the case $d = 2$, both Theorems 3.1 and 3.6 can be proved more directly, exploiting the fact that nonparallel lines in $\mathbb{R}^2$ always meet. The resulting Π-paths are then formed from consecutive sequences of line segments drawn from Π. However our interest is in Π-*geodesics*, and even in case $d = 2$ it is not yet known whether Π-geodesics can be constructed as consecutive sequences of line segments.

While we define Π-geodesics as minimum-*time* paths, we retain an interest in the actual lengths of Π-geodesics. It is a consequence of Theorem 2.6 that a Π-geodesic between two points is almost surely of finite length: this follows because if the Π-geodesic has finite duration T then it must be contained in a sufficiently large ball, and therefore its maximum speed is bounded, which in turn bounds the length. There is a more subtle question, namely whether the *mean* length of the Π-geodesic is finite. We shall answer this question in the affirmative in Section 5, but only for the case of dimension $d = 2$.

The above results establish the existence of Π-geodesics, but only nonconstructively. The principal difficulty in taking a constructive approach lies in the implicit tree-like way in which Π-paths are constructed in Theorems 3.1 and 3.6. In the remainder of this section, we show how to approximate Π-paths by sequentially-defined Lipschitz paths which are almost Π-paths. The major benefit of this result is that it implies the measurability of the random time which a Π-geodesic would take to move from one specified point x to another specified point y. As is commonly the case for measurability arguments, the details are a little tedious; however, the result does provide theoretical justification for some simulation constructions of Π-geodesics (e.g., the construction in Figure 1). The essence of the

matter is to work with Lipschitz paths which would be Π-paths if the upper-semicontinuous speed-limit V were replaced by $\max\{\varepsilon, V\}$ for some small $\varepsilon > 0$. The methods of proof of the following results also justify the simulation algorithm used to produce the realizations of networks in Figure 1.

DEFINITION 3.7 (ε-near-sequential-Π-path). For given $\varepsilon > 0$, a continuous path $\widetilde{\xi} : [0, T] \to \mathbb{R}^d$ is an *ε-near-sequential-Π-path* if there is a finite dissection of the interval $[0, T]$ as

$$0 = b_0 \le a_1 \le b_1 \le a_2 \le b_2 \le \cdots \le a_m \le b_m \le a_{m+1} = T,$$

associated with a finite sequence of marked lines from Π (possibly with repeats),

$$(\widetilde{\ell}_1, \widetilde{v}_1), (\widetilde{\ell}_2, \widetilde{v}_2), \ldots, (\widetilde{\ell}_m, \widetilde{v}_m),$$

such that:

(a) $\widetilde{\xi}(t) \in \widetilde{\ell}_r$ when $a_r \le t \le b_r$, and $|\widetilde{\xi}'(t)| \le \widetilde{v}_r$ for almost all $t \in (a_r, b_r)$;
(b) $|\widetilde{\xi}'(t)| < \varepsilon$ for almost all $t \in \bigcup_{r=0}^{m} [b_r, a_{r+1}]$ and $\widetilde{\xi}'$ is constant on each $[b_r, a_{r+1}]$;
(c) $\sum_{r=0}^{m} |a_{r+1} - b_r| < \varepsilon$.

The next result shows that Π-paths can be approximated by ε-near-sequential-Π-paths for small $\varepsilon > 0$, simply by using the principal marked lines involved in ξ to generate the finite marked line sequence $(\widetilde{\ell}_1, \widetilde{v}_1), (\widetilde{\ell}_2, \widetilde{v}_2), \ldots, (\widetilde{\ell}_m, \widetilde{v}_m)$.

THEOREM 3.8. *Suppose only that $\gamma > 1$, so that the line processes Π_v are locally finite for each $v > 0$. Consider a Π-path $\xi : [0, T] \to \mathbb{R}^d$, defined up to some fixed finite time T and running from x_1 to x_2. For each $\varepsilon > 0$ there can be found an ε-near-sequential-Π-path $\widetilde{\xi} : [0, T] \to \mathbb{R}^d$ such that*:

$\widetilde{\xi}(0) = x_1, \widetilde{\xi}(T) = x_2$;
$\sup\{|\xi(t) - \widetilde{\xi}(t)| : t \in [0, T]\} < \varepsilon$.

Before proving this, we state the following important corollary, whose proof is an immediate consequence.

COROLLARY 3.9. *Suppose only that $\gamma > 1$. Every Π-path defined up to finite time T can be uniformly approximated by a sequence of ε-near-sequential-Π-paths such that $\varepsilon \downarrow 0$ along the sequence.*

PROOF OF THEOREM 3.8. It suffices to prove the result for a fixed positive $\varepsilon < 1$.

The Poisson line process Π has no triple intersections and, therefore, a given Π-path ξ can only ever lie on at most two lines simultaneously. Hence, by countable exhaustion (based on ordering by $\mathrm{Leb}\{t \in [0, T] : \xi(t) \in \ell\}$) it follows that

there can be only countably many marked lines $(\ell, v) \in \Pi$ such that $\mathrm{Leb}\{t \in [0,T] : \xi(t) \in \ell\} > 0$ (so that ℓ and ξ intersect in a time-set of positive measure).

Since ξ is continuous, we know that the image $\mathrm{Im}(\xi)$ of ξ is compact and, therefore (since Π_v is locally finite for any positive v), the lines intersecting ξ in time-sets of positive measure can be sequentially ordered by speed in decreasing order:

$$(\ell_1, v_1), (\ell_2, v_2), (\ell_3, v_3), \dots \quad \text{and} \quad v_1 \geq v_2 \geq v_3 \geq \cdots. \tag{3.3}$$

Note that we do *not* presume that the ℓ_i are visited sequentially in order.

It is an immediate consequence of the above that ξ is $\mathrm{Lip}(v_1)$. Moreover, a consequence of the countable exhaustion construction is that

$$\mathrm{Leb}\{t \in [0,T] : \xi(t) \notin \ell_1 \cup \ell_2 \cup \cdots \cup \ell_n\} \to 0 \qquad \text{as } n \to \infty.$$

Therefore, for all $\varepsilon > 0$, for all sufficiently large n,

$$\mathrm{Leb}\{t \in [0,T] : \xi(t) \notin \ell_1 \cup \ell_2 \cup \cdots \cup \ell_n\} < \varepsilon \quad \text{and} \quad v_{n+1} < \varepsilon.$$

We use the finite sequence $(\ell_1, v_1), (\ell_2, v_2), \dots, (\ell_n, v_n)$ to generate an ε-near-sequential-Π-path $\widetilde{\xi}$ which approximates ξ in uniform norm. Note that this finite sequence is *not* the same as the finite sequence $(\widetilde{\ell}_1, \widetilde{v}_1), (\widetilde{\ell}_2, \widetilde{v}_2), \dots, (\widetilde{\ell}_m, \widetilde{v}_m)$ from Definition 3.7, but is used to construct it iteratively.

We begin by setting $\widetilde{\xi}(0) = \xi(0)$ and $\widetilde{\xi}(T) = \xi(T)$.

Consider first the time set $\{t \in [0,T] : \xi(t) \in \ell_1\}$. The Π-path ξ makes a countable number of excursions away from ℓ, and the excursion intervals form the connected components of the (relatively) open set $[0,T] \setminus \{t \in [0,T] : \xi(t) \in \ell_1\}$. There are at most two incomplete excursions in the interval $[0,T]$, namely the beginning excursion, anchored to x_1 at time 0 on the left, and the end excursion, anchored to x_2 at time T on the right. In addition, there can be at most finitely many complete excursions for which $\mathrm{dist}(\xi, \ell)$ reaches the level ε [in fact, a calculation using the $\mathrm{Lip}(v_1)$ property of ξ gives an upper bound on the number of such excursions, namely $\frac{1}{2} v_1 T/\varepsilon$]. We set

$$U_0 = [0,T],$$

$$U_1 = \{t \in U_0 : \xi(t) \notin \ell_1\} \setminus \bigcup \{(a,b) \subset U_0 : \xi(a), \xi(b) \in \ell_1 \text{ and } 0 < \mathrm{dist}(\xi(t), \ell_1) < \varepsilon \text{ for } a < t < b\}.$$

So U_1 is a finite union of intervals (relatively open in U_0).

Define $\widetilde{\xi}$ on $U_0 \setminus U_1$ as the orthogonal projection of ξ on ℓ_1. By the properties of orthogonal projection and the $\mathrm{Lip}(v_1)$ property of ξ on U_0, it follows that $|\widetilde{\xi}'(t)| \leq v_1$ for almost all $t \in U_0 \setminus U_1$. Moreover, by construction,

$$|\widetilde{\xi}(t) - \xi(t)| = 0 \qquad \text{if } \xi(t) \in \ell_1 \text{ and } t \in U_0 \setminus U_1,$$

$$|\widetilde{\xi}(t) - \xi(t)| < \varepsilon \qquad \text{for other } t \in U_0 \setminus U_1.$$

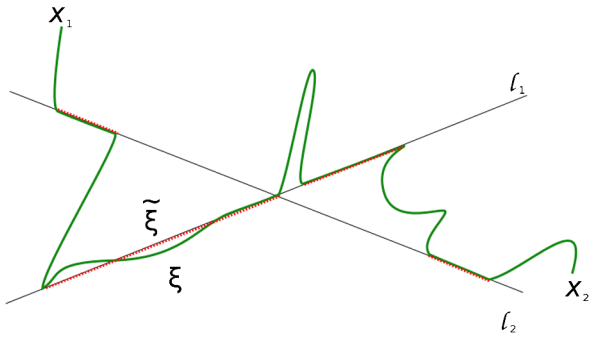

FIG. 5. *First two stages of the iterative construction of an ε-near-sequential-Π-path. The solid curve represents the trajectory of the Π-path ξ. The dotted segments represent the partially-defined trajectory of $\widetilde{\xi}$ after these first two stages (later stages successively fill in the gaps). Note that $\widetilde{\xi}$ is defined* (a) *when ξ runs along one of the lines ℓ_1, ℓ_2 and* (b) *when ξ makes small excursions from one of these lines.*

Hence, $|\widetilde{\xi}(t) - \xi(t)| < \varepsilon$ for all $t \in U_0 \setminus U_1$. Finally, note that $\{t \in [0,T] : \xi(t) \in \ell_1\} \subseteq [0,T] \setminus U_1$.

Now define

$$U_2 = \{t \in U_1 : \xi(t) \notin \ell_2\} \setminus \bigcup \{(a,b) \subset U_1 : \xi(a), \xi(b) \in \ell_2 \text{ and } 0 < \operatorname{dist}(\xi(t), \ell_2) < \varepsilon \text{ for } a < t < b\},$$

and note that by construction ℓ_1 cannot intersect ξ in a time-set of positive measure in U_1, so that ξ is $\operatorname{Lip}(v_2)$ in the time-set U_1. We can argue as before that U_2 is a finite union of relatively open intervals. We can extend the definition of $\widetilde{\xi}$ to $U_1 \setminus U_2$ by using orthogonal projection of ξ onto ℓ_2: we have $|\widetilde{\xi}'(t)| \le v_2$ for almost all $t \in U_1 \setminus U_2$, and $|\widetilde{\xi}(t) - \xi(t)| < \varepsilon$ for all $t \in U_1 \setminus U_2$. Moreover, $\{t \in [0,T] : \xi(t) \in \ell_1 \cup \ell_2\} \subseteq [0,T] \setminus U_2$.

The construction is illustrated in Figure 5.

Iterating this construction, we end up defining $\widetilde{\xi}$ on a time-set $[0,T] \setminus U_n$ containing $\{t \in [0,T] : \xi(t) \in \ell_1 \cup \cdots \cup \ell_n\}$, such that $\widetilde{\xi}(t) \in \ell_1 \cup \ell_2 \cup \cdots \cup \ell_n$ for $t \in [0,T] \setminus U_n$, with $|\widetilde{\xi}'(t)| \le v_r$ for almost all t such that $\widetilde{\xi}(t) \in \ell_r$, for $r = 1, 2, \ldots, n$, and finally $|\widetilde{\xi}(t) - \xi(t)| < \varepsilon$ if $t \in [0,T] \setminus U_n$. Since $\{t \in [0,T] : \xi(t) \in \ell_1 \cup \cdots \cup \ell_n\} \subseteq [0,T] \setminus U_n$, it follows from the countable exhaustion construction that $\operatorname{Leb}([0,T] \setminus U_n) < \varepsilon$.

We next complete the construction on the finite family of excursion intervals which are the connected components of the relatively open set U_n. Note that by construction $\widetilde{\xi}$ agrees with ξ on the end-points of these excursion intervals. None of the lines $\ell_1, \ell_2, \ldots, \ell_n$ intersect U_n in a time-set of positive measure: therefore ξ satisfies a $\operatorname{Lip}(v_{n+1})$ property on U_n. Hence, for each of these excursion intervals, if a and b are the end-points then $|\xi(b) - \xi(a)| \le (b-a)v_{n+1} < (b-a)\varepsilon$.

Accordingly, we can define $\widetilde{\xi}$ by linear interpolation over the excursion interval (so that $\widetilde{\xi}'$ is indeed constant over this excursion interval), with the result that

$|\widetilde{\xi}'(t)| < \varepsilon$ for almost all $t \in [a, b]$. Finally, the Lip(v_{n+1}) property implies that $|\xi(t) - \xi(a)|$ and $|\xi(t) - \xi(b)|$ are both strictly bounded above by $|b-a|\varepsilon \leq \varepsilon^2 \leq \varepsilon$ when $a < t < b$ (use $|b - a| \leq \text{Leb}([0, T] \setminus U_n) < \varepsilon$ and $\varepsilon \leq 1$); it follows by convexity that the same bound holds if a and b are replaced by the piecewise interpolant $\widetilde{\xi}(t)$:

$$\left|\xi(t) - \widetilde{\xi}(t)\right| < \varepsilon.$$

It follows that $\widetilde{\xi}$ is the required ε-near-sequential-Π-path approximating ξ to within ε in uniform norm. The sequence $(\widetilde{\ell}_1, \widetilde{v}_1), (\widetilde{\ell}_2, \widetilde{v}_2), \ldots, (\widetilde{\ell}_m, \widetilde{v}_m)$ is obtained from the successive visits (with repetitions) of $\widetilde{\xi}$ to the finite sequence of lines $(\ell_1, v_1), (\ell_2, v_2), \ldots, (\ell_n, v_n)$. □

REMARK 3.10. If ξ is a Π-geodesic, then the above construction can be simplified. Using the notation of the proof, suppose that (a, b) is a connected component of U_n. The maximum speed of ξ in U_n is v_{n+1}: consequently, if $\xi(s)$ and $\xi(t)$ belong to ℓ_{n+1} for $s < t$, both belonging to (a, b), then the fastest route from $\xi(s)$ to $\xi(t)$ must run along ℓ_{n+1} at maximum permitted speed v_n. Consequently, $\{t \in (a, b) : \xi \in \ell_{n+1}\}$ must already be a relatively closed interval in (a, b), so that there is no need to use the excursion construction in the proof of the theorem.

We have seen that, under the weak condition $\gamma > 1$, every Π-path can be uniformly approximated by a sequence of ε-near-sequential-Π-paths with $\varepsilon \downarrow 0$. Conversely, if we strengthen the condition on γ to $\gamma \geq d$ (so that the a priori bound of Theorem 2.6 is available) then there is a kind of compactness result for ε-near-sequential-Π-paths.

THEOREM 3.11. *Suppose that $\gamma \geq d \geq 2$ and $T < \infty$. For $n = 1, 2, \ldots$, let $\widetilde{\xi}_n : [0, T] \to \mathbb{R}^d$ be an ε_n-near-sequential-Π-path from x to y, with $\varepsilon_n \downarrow 0$. Then there are subsequences $\{\widetilde{\xi}_{n_k} : k = 1, 2, \ldots\}$ which converge uniformly to Π-path limits.*

PROOF. Since ε_n is decreasing in n, each ε_n-near-sequential-Π-path $\widetilde{\xi}_n$ obeys the single modified speed limit $\max\{\varepsilon_1, V\}$. Hence, the comparison argument of Theorem 2.6 can be adapted to show that all the ε_n-near-sequential-Π-paths $\widetilde{\xi}_1$, $\widetilde{\xi}_2, \ldots$ lie in a single ball B of radius R depending on V and ε_1.

Consequently, the $\widetilde{\xi}_1$, $\widetilde{\xi}_2, \ldots$ obey a uniform Lipschitz condition (with Lipschitz constant given by the speed of the fastest line to hit the ball B); therefore by the Arzela–Ascoli theorem we can extract a uniformly convergent subsequence $\{\widetilde{\xi}_{n_k} : k = 1, 2, \ldots\}$ whose limit $\widetilde{\xi}_\infty$ is also a Lipschitz path with the same Lipschitz constant.

The persistence of Lipschitz constants in the limit also holds locally. For fixed $\lambda > 0$, consider the open set $\mathcal{S}_v^c = \{x : V(x) < v\}$. Fix $0 < s < t < T$ such that

$\mathrm{Im}(\widetilde{\xi}_\infty|_{[s,t]}) \subset \mathcal{S}_v^c$. But $\widetilde{\xi}_\infty$ is continuous, so $\mathrm{Im}(\widetilde{\xi}_\infty|_{[s,t]})$ is compact; therefore the uniform convergence of $\widetilde{\xi}_{n_k} \to \widetilde{\xi}_\infty$ implies that for all $k \geq k_\lambda$ we have

$$\mathrm{Im}(\widetilde{\xi}_n|_{[s,t]}) \subset \mathcal{S}_v^c.$$

It follows that if $k \geq k_\lambda$ then $\widetilde{\xi}_{n_k}$ satisfies a $\mathrm{Lip}(\max\{\varepsilon_{n_k}, v\})$ condition over the time set $[s,t]$. Bearing in mind that $\varepsilon_n \downarrow 0$, we deduce that $\widetilde{\xi}_\infty$ satisfies a $\mathrm{Lip}(v)$ condition whenever $\widetilde{\xi}_\infty$ belongs to $\mathcal{S}_v^c$. This implies that the following is a Lebesgue-null subset of $[0, T]$:

$$\{t \in [0,T] : |\widetilde{\xi}'_\infty(t)| > v \text{ and } \widetilde{\xi}_\infty(t) \notin \mathcal{S}_v\}.$$

Thus, the subsequential limit $\widetilde{\xi}_\infty$ is a Π-path (Definition 2.3). □

This allows us to deduce the measurability of the random variable which is given by the time taken for a Π-geodesic to move between specified end-points x_1 and x_2.

COROLLARY 3.12. *Suppose that $\gamma > d$. Fix x_1 and x_2 in $\mathbb{R}^d$, and let T be the least time such that there is a Π-path running from x_1 to x_2 in time T, equivalently, such that the (possibly nonunique) Π-geodesic from x_1 to x_2 has duration T. Then T is a function of the marked line process Π: it is in fact measurable, and hence a random variable.*

PROOF. Consider the event $E_{\varepsilon,\tau}$ that there is an ε-near-sequential-Π-path from x_1 to x_2 of duration at most τ. This event is measurable, because we may restrict attention to a countable subfamily of ε-near-sequential-Π-paths, determined, for example, so that constituent line segments are bounded by the intersection point process of Π.

But it is a consequence of the above results that

$$\bigcap_\varepsilon E_{\varepsilon,\tau} = [\text{duration of } \Pi\text{-geodesic from } x_1 \text{ to } x_2 \text{ is no more than } \tau].$$

For Theorem 3.8 shows that the existence of such a Π-geodesic leads to the construction of ε-near-sequential-Π-paths from x_1 to x_2 of the same duration as the Π-geodesic. On the other hand, it is an immediate consequence of Theorem 3.11 that if $E_{\varepsilon_n,\tau}$ is nonempty for a sequence $\varepsilon_n \downarrow 0$ then there must exist a Π-path from x_1 to x_2 of duration τ. (Note that we may prolong the duration of any ε-near-sequential-Π-path simply by holding it at its destination.) Finally, the events $E_{\varepsilon,\tau}$ are decreasing in n, so the intersection $\bigcap_\varepsilon E_{\varepsilon,\tau}$ can be reduced to a countable intersection. It follows that

$$[\text{duration of } \Pi\text{-geodesic from } x_1 \text{ to } x_2 \text{ is no more than } \tau]$$

is a measurable event. □

We note that simple selection criteria can be used to establish measurable maps which yield intervening Π-geodesics for each pair of end-points x_1 and x_2.

4. Π-geodesics: Almost-sure uniqueness in dimension 2. In the simplest nontrivial case, namely $d = 2$, it can be shown that the Π-geodesic between two specified points is almost surely unique. The method of proof makes essential use of the point-line duality of the plane, so does not extend to the case $d \geq 3$.

REMARK 4.1. The assertion of almost sure uniqueness between of Π-geodesic connections between two specified points does *not* imply that almost surely all pairs of points are connected by unique Π-geodesics: a simple counterexample can be constructed by considering the possibility that three lines, of speeds only just below unit speed, form an approximate equilateral triangle Δ of near unit-side length. Let ρ be the perimeter of Δ, and suppose all other lines within $\rho/4$ of Δ are of less than speed $1/2$, while all lines hitting the interior of Δ are of speed substantially less than $1/2$. (How much less depends on the approximation to equilateral shape.) Both these events have positive probability. Then consider the two Π-paths of length $\rho/2$ running either way round Δ from a given reference point on the boundary of Δ. These form two distinct Π-geodesics between the same end-points. (The construction is illustrated in Figure 6.)

Note that structures similar to this counterexample will exist at all scales in the random metric space produced from $\mathbb{R}^2$ furnished with a random metric derived

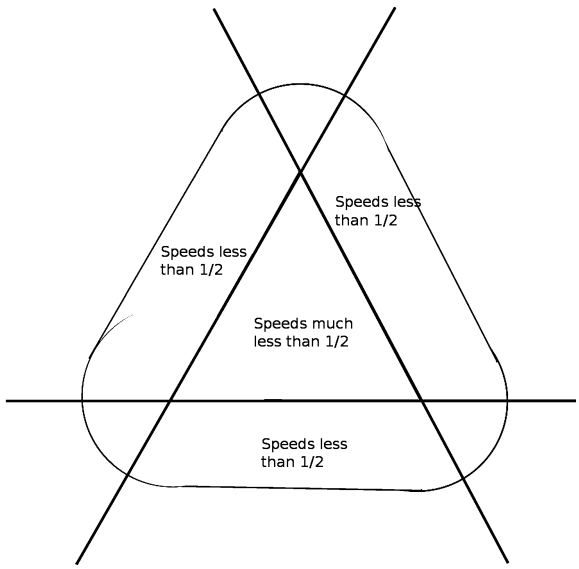

FIG. 6. *Within the approximate equilateral triangle delineated by three fast lines, speeds are slow enough to prevent short-cuts. Outside the triangle, up to a distance of one quarter of the perimeter, speeds are slow enough that there is no possibility of* Π*-geodesics looping outside this region and then returning to the triangle.*

from $\Pi = \Pi^{(2,\gamma)}$ for $\gamma > 2$. So these random metric spaces are far from hyperbolic in the sense of Gromov (defined, e.g., in Burago, Burago and Ivanov [12], Section 8.4).

We begin with a structural result about Π-geodesics in dimension 2, namely that if ℓ is a line from Π (hence of positive speed limit) which contributes a segment to a Π-geodesic ξ then ξ joins and leaves ℓ using simple intersections of ℓ with other lines in Π.

DEFINITION 4.2 (Proper encounter of a line by a Π-path). Suppose $d = 2$ and $\gamma > 2$. We say that a π-path ξ *encounters a line* ℓ *of* Π *properly* at $\xi(t) \in \ell$ if there is $\varepsilon > 0$ such that within ball$(\xi(t), \varepsilon)$ the Π-path ξ is not contained solely in ℓ, but lies in the union of ℓ and a further line $\widetilde{\ell}$ from Π.

The notion of a proper encounter is vacuous for Π-paths in the case $d > 2$, because then almost surely lines of the Poisson line process Π do not intersect each other.

THEOREM 4.3. *Suppose $d = 2$ and $\gamma > 2$. With probability* 1, *for each line $\ell \in \Pi$ and each Π-geodesic ξ, the intersections of ξ with ℓ form a finite disjoint union of intervals, such that the nonsingleton intervals are delimited by proper encounters of ξ with ℓ.*

PROOF. First, note that with probability 1 there are no triple intersections of lines ℓ_1, ℓ_2, ℓ_3 from Π. Given this, the remainder of the argument is purely geometric and, therefore, applies to all Π-geodesics simultaneously.

Consider the set of all lines ℓ in Π intersecting ξ. As noted in Remark 3.10, the Π-geodesic property implies that the intersection of ξ with the fastest such line must be a single (possibly trivial) interval. This is because the fastest route between first and last intersection with this fastest line must lie along the line. For the second fastest line, the intersection must be formed as the union of at most two intervals, which must be encountered, respectively, before and after the encounter with the fastest line. Continuing this argument, the intersection of ξ with the kth fastest line must be the union of at most 2^{k-1} intervals. Thus, for *any* line ℓ from Π, if ξ intersects ℓ at all then the intersection set must be the union of a finite number of intervals, some of which may be trivial.

For a given Π-geodesic ξ, fix a given line ℓ_1 from Π which intersects ξ, and consider the start $\xi(t)$ of a nonsingleton intersection interval $[\xi(t), \xi(s)]$ between ξ and ℓ_1. [Reversing time, the following argument will apply to the departure point $\xi(s)$ as well as to the arrival point $\xi(t)$.] For sufficiently small ε, either ℓ_1 will be the fastest line in ball$(\xi(t), \varepsilon)$, or it will be the second fastest, and the fastest line ℓ_2 will intersect ℓ_1 at $\xi(t)$. Choose $u < t$ to be as small as possible subject to the requirement that $\xi|_{(u,t]} \subset \text{ball}(\xi(t), \varepsilon)$. It follows from the Π-geodesic property that if $\xi|_{(u,t)}$ hits ℓ_2 at time $v \in (u, t)$ (assuming ℓ_2 exists) then $\xi|_{[v,t]}$ must

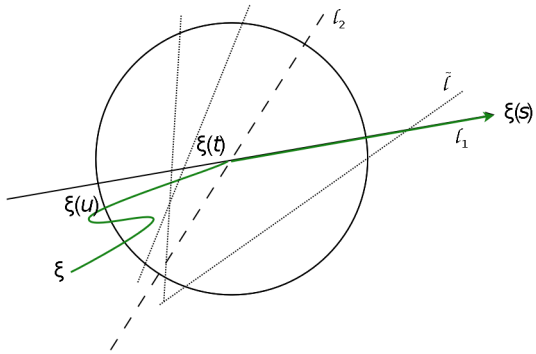

FIG. 7. *Illustration of the construction used in the proof of Theorem 4.3 to demonstrate that a geodesic ξ must make a proper encounter on a line $\ell_1 \in \Pi$. Here, ℓ_2 is a possible faster line; by making the enclosing ball small enough we can exclude the possibility that such a faster ℓ_2 hits ℓ_1 in any place other than the point $\xi(t)$ where ξ hits ℓ_1 for the first time. Dotted lines are lines of low cost, where the notion of cost is described in the construction.*

run along ℓ_2, so that ξ makes a proper encounter with ℓ_1 using the line ℓ_2. On the other hand, if $\xi|_{(u,t)}$ does not hit ℓ_2 then it cannot hit ℓ_1, since if it did so at time $v \in (u,t)$ then $\xi|_{[v,t]}$ would have to run along ℓ_1, contradicting the maximality of $[t,s]$. Thus, we may restrict attention to the case when $\xi|_{(u,t)}$ hits neither ℓ_1 nor ℓ_2. This and following features of the construction are illustrated in Figure 7.

We introduce the notion of *cost*, based on comparison between the motion defined by ξ over (u,t) (say) and the motion defined by a comparison particle $\widetilde{\xi}$, which begins at time u at the location which is the projection of $\xi(u)$ on ℓ_1, and which continues at maximum speed (w, say) along ℓ_1 in the direction from $\xi(t)$ to $\xi(s)$. We compute the cost of following ξ rather than $\widetilde{\xi}$ in terms of the time by which $\widetilde{\xi}$ leads ξ when ξ hits ℓ_1 (namely, at time t). This is given by the integral

$$\frac{1}{w}\int_u^t \bigl(w - |\xi'(a)|\cos\theta(a)\bigr)\,\mathrm{d}a = \frac{1}{w}\int_u^t \bigl(w - v(a)\cos\theta(a)\bigr)\,\mathrm{d}a,$$

where $\theta(a)$ is the angle that $\xi'(a)$ makes with ℓ, and setting $v(a) = |\xi'(a)|$. [Note that $v(a) = |\xi'(a)| = V(\xi(a))$, because ξ is a Π-geodesic.]

Let H be the perpendicular distance between $\xi(u)$ and ℓ_1, and then reparameterize the above equation in terms of the perpendicular distance between $\xi(a)$ and ℓ_1, removing those parts of the integral for which $\xi'(a)$ is not directed toward ℓ_1 [in which case the contribution to the integral is certainly positive, since ℓ_1 is faster than any other line used by ξ over (u,t)]. Setting

$$\overline{a}(h) = \inf\bigl\{a\colon \text{ perpendicular distance of } \xi(a) \text{ from } \ell_1 \text{ is } h\bigr\},$$

and $\overline{v}(h) = v(\overline{a}(h))$, and $\overline{\theta}(h) = \theta(\overline{a}(h))$, we find

$$\frac{1}{w}\int_u^t \big(w - |\xi'(a)|\cos\theta(a)\big)\,\mathrm{d}a \geq \frac{1}{w}\int_0^H \big(w - \overline{v}(h)\cos\overline{\theta}(h)\big)\frac{\mathrm{d}h}{\overline{v}(h)\sin\overline{\theta}(h)}$$
$$= \int_0^H \left(\frac{\csc\overline{\theta}}{\overline{v}} - \frac{\cot\overline{\theta}}{w}\right)\mathrm{d}h.$$

Accordingly, define the relative *cost index* of a given line ℓ from Π (compared with ℓ_1) by

$$c(\ell) = \frac{\csc\theta}{v} - \frac{\cot\theta}{w}, \tag{4.1}$$

where v is the speed limit of ℓ, and θ is the angle it makes with ℓ_1. Evidently, the time by which $\widetilde{\xi}$ leads ξ when ξ hits ℓ_1 can be controlled in terms of an integral of cost indices of lines along which ξ travels when directed toward ℓ_1. The cost index of ℓ_1 is not defined, though a limiting argument gives the value 0. Note too that, for any line ℓ of speed v, the cost index of ℓ turns out to be positive if $v < w$.

Consider line-space parametrized using $\xi(t)$ and ℓ_1 as reference point and reference line, restricting attention to lines with speed limit less than w (the speed limit for ℓ_1). The intensity measure $\frac{\gamma-1}{2}v^{-\gamma}\,\mathrm{d}v\,\mathrm{d}r\,\mathrm{d}\theta$ may be re-expressed in terms of c, r and θ: since

$$\frac{\mathrm{d}c}{\mathrm{d}v} = -\frac{\csc\theta}{v^2},$$

it follows that in the new coordinates the intensity measure is

$$\frac{\gamma-1}{2}\sin\theta\left(c\sin\theta + \frac{\cos\theta}{w}\right)^{\gamma-2}\mathrm{d}c\,\mathrm{d}r\,\mathrm{d}\theta. \tag{4.2}$$

Now $v^{-1} = c\sin\theta + \frac{\cos\theta}{w} > 0$, so the measure determined by (4.2) is nonnegative. Consider the line pattern of lines with cost smaller than a specified threshold c_0. From the form of (4.2), this pattern is locally finite. On the other hand, the line pattern of lines with cost exceeding a specified threshold must be locally dense even when constrained by requiring angle θ to lie within a small interval.

We pick $\widetilde{\ell}$ to be the lowest cost line separating $\xi(s)$ from $\xi_{(u,t)}$ (see Figure 7 for a possible configuration, notice that this line may or may not be ℓ_2), and we determine the minimal $v \in [u,t]$ such that all the lines involved in $\xi|_{(v,t)}$ are more expensive than $\widetilde{\ell}$.

If $v = t$, then consider the lim inf as $\delta \to 0$ of the costs of lines involved in $\xi|_{(t-\delta,t)}$. This must be finite, for otherwise a low-cost line would produce a path faster than the Π-geodesic. Since there are only finitely many low-cost lines near $\xi(t)$, this means there must be at least one low-cost line which is repeatedly visited by ξ in every interval $(t-\delta, t)$; so this line must pass through $\xi(t)$ and, therefore, must either be ℓ_1 (which we have excluded) or ℓ_2 (which is a case already disposed of). Hence, we may suppose $v < t$.

If $v < t$, pick the line ℓ^* of least cost which is hit by $\xi|_{(v,t)}$. Suppose (a) this hits the component of $\ell_1 \setminus \xi(s)$ not containing $\xi(t)$. Then a combination of ℓ^* and $\widetilde{\ell}$ and ℓ_1 provides a faster way to get to $\xi(s)$ than is provided by ξ, again violating the Π-geodesic property of ξ. Otherwise, (b) this line ℓ^* does not hit the component of $\ell_1 \setminus \xi(s)$ not containing $\xi(t)$. If ξ does not meet ℓ_1 at $\xi(t)$ using ℓ^*, then ℓ^* followed by ℓ_1 provides a faster way to get to $\xi(s)$ than is provided by ξ, violating the Π-geodesic property of ξ.

It follows that if ξ does not meet ℓ_1 at $\xi(t)$ using ℓ_2 then it must meet ℓ_1 at $\xi(t)$ using ℓ^*, proving the result.

As noted above, a time-reversal argument deals with the departure time s. Consequently, all the countably many nonsingleton intersections of ξ with lines of positive speed limit must be proper. □

Note in passing that in higher dimension $d > 2$ the quantity analogous to the cost (4.1) varies along each line.

We can now prove almost sure uniqueness of Π-geodesics between specified pairs of points in the planar case.

THEOREM 4.4. *Suppose $\gamma > d = 2$. Consider two points x and y in the plane $\mathbb{R}^2$. Almost surely there is just one Π-geodesic connecting x and y.*

PROOF. First, note the following consequence of Theorem 2.6: as $R \to \infty$, so

$$\mathbb{P}\big[\text{all geodesics from } x \text{ to } y \text{ are contained in ball}(\mathbf{o}, R)\big] \to 1.$$

Furthermore, Theorem 4.3 implies that the following assertion holds almost surely: all Π-geodesics join or leave any line ℓ in Π at only countably many possible places, namely the intersection points of ℓ with the rest of Π. Moreover, any particular Π-geodesic joins or leaves any particular line at only finitely many of these places.

Finally, note that if two Π-geodesics from x to y intersect at z then they must do so at the same relative time as measured from x.

Bearing these facts in mind, we now develop the proof.

For fixed $v > 0$, pick a line ℓ_0 uniformly at random from Π_{v_0} such that $\ell_0 \Uparrow \text{ball}(\mathbf{o}, R)$. Note that the speed V of ℓ_0 has a Pareto distribution, with density $(\gamma - 1)(v/v_0)^{-\gamma}$ for $v > v_0$. Note too that V is independent of the physical location of ℓ_0 and (by Slivnyak's theorem) is independent of $\Pi \setminus \{\ell_0\}$ which itself is a statistical copy of Π. Then (almost surely) for all sufficiently large $R > 0$ we know all Π-geodesics from x to y belong to the path-set $\mathcal{P}^{\ell_0}(x, y) = \bigcup_{u>0} \mathcal{P}_u^{\ell_0}(x, y)$, where $\mathcal{P}_u^{\ell_0}(x, y)$ is the set of Π-paths from x to y lying in ball$(\mathbf{o}, R)$ for which:

- the Π-path is always run at maximal permissible speed;
- excursions of the Π-path away from ℓ_0 are Π-geodesics;

- the intersections of the Π-path with ℓ_0 form a finite disjoint union of intervals $[a, b]$;
- and from these intervals the nonsingleton intervals are delineated by proper encounters of the Π-path and ℓ_0, moreover these end-points lie in $\ell_0 \cap \bigcup\{\ell : \ell \in \Pi_u\}$.

We further decompose $\mathcal{P}_u^{\ell_0}(x, y) = \bigcup_\sigma \mathcal{P}_u^{\ell_0}(x, y; \sigma)$, where σ ranges over the family of finite sequences of pairs of (signed) integers, such that the closed intervals delineated by the pairs of integers are disjoint. To define $\mathcal{P}_u^{\ell_0}(x, y; \sigma)$, consider the doubly-infinite point sequence $\ell \cap \bigcup\{\ell : \ell \in \Pi_u\}$, and index the points by $\mathbb{Z}$ once and for all, using a fixed sense of direction along ℓ_0 and arranging for the interval determined by the points indexed by 0 and 1 to be the (almost surely unique) interval nearest to $\mathbf{o}$. Then $\mathcal{P}_u^{\ell_0}(x, y; \sigma)$ is composed of those Π-paths in $\mathcal{P}_u(x, y)$ for which the union of disjoint nonsingleton intervals of intersection with ℓ_0 equals the union of the intervals bounded by pairs of points indexed by the pairs of σ, moreover σ lists these intervals in the order in which they are visited and according to the direction in which they are travelled.

It is a consequence of the defining properties of $\mathcal{P}^{\ell_0}(x, y)$ etc., and the property that intersecting Π-geodesics from x to y must visit their intersections at the same relative times, that all Π-paths in $\mathcal{P}_u^{\ell_0}(x, y; \sigma)$ spend the same amount of time S_σ outside ℓ_0. Moreover, consider the lengths $L_{\sigma_1}, L_{\sigma_2}, \ldots, L_{\sigma_k}$ corresponding to the intervals bounded by pairs of points indexed by the elements of $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_k)$. These are sums of independent Gamma random variables of the same rate, and all Π-paths in $\mathcal{P}_u^{\ell_0}(x, y; \sigma)$ spend the same amount of time $L_\sigma / V = (L_{\sigma_1} + L_{\sigma_2} + \cdots + L_{\sigma_k})/V$ on ℓ_0. Moreover, the S_σ and the L_{σ_i} random variables are statistically independent of the speed V of ℓ_0.

If $\sigma \neq \widetilde{\sigma}$, then the sums $L_\sigma = L_{\sigma_1} + L_{\sigma_2} + \cdots + L_{\sigma_k}$, $L_{\widetilde{\sigma}} = L_{\widetilde{\sigma}_1} + L_{\widetilde{\sigma}_2} + \cdots + L_{\widetilde{\sigma}_{\widetilde{k}}}$ can be decomposed into summands over shared or distinct Gamma random variables to reveal that $L_\sigma - L_{\widetilde{\sigma}}$ has a nondegenerate probability density whenever $\sigma \neq \widetilde{\sigma}$ and, therefore, $\mathbb{P}[L_\sigma = L_{\widetilde{\sigma}}] = 0$. Hence, Π-paths from $\mathcal{P}_u^{\ell_0}(x, y; \sigma)$ have a common duration of $S_\sigma + L_\sigma / V$, and Π-paths from $\mathcal{P}_u^{\ell_0}(x, y; \widetilde{\sigma})$ have a common duration of $S_{\widetilde{\sigma}} + L_{\widetilde{\sigma}} / V$, and the Pareto distribution and independence of V implies that

$$\mathbb{P}[S_\sigma + L_\sigma / V = S_{\widetilde{\sigma}} + L_{\widetilde{\sigma}} / V] = 0 \qquad \text{if } \sigma \neq \widetilde{\sigma}.$$

Thus, almost surely, for all the countably many different $\sigma \neq \widetilde{\sigma}$, the common durations of Π-paths from $\mathcal{P}_u^{\ell_0}(x, y; \sigma)$ and $\mathcal{P}_u^{\ell_0}(x, y; \widetilde{\sigma})$ are different.

It follows that almost surely all the Π-geodesics between x and y must traverse the same set of nonsingleton intervals in the same direction along ℓ_0, since any two such Π-geodesics will have to belong to the same $\mathcal{P}_u^{\ell_0}(x, y; \sigma)$ for sufficiently small $u > 0$. But this must then hold for all $\ell \in \Pi$, and therefore (since

Π-geodesics must intersect at the same time as measured from x) almost surely all Π-geodesics between x and y must agree. $\square$

The argument here is delicate: for example, it is *not* the case that the set of lengths along lines between intersections is linearly independent if we consider the ensemble of lengths of a unit Poisson line process. Indeed, the tessellation produced by a Poisson line process will be rigid; consideration of various triangles shows that the length of any single segment will be determined by the lengths of all the other segments, so long as the incidence geometry of the segments is known.

The almost-sure uniqueness of planar Π-geodesics implies that planar spatial networks formed from the Poisson line process model satisfy property 1 of Definition 1.1.

5. Π-geodesics: Finiteness of mean-length in dimension 2. One might conceive that a Π-geodesic between two fixed points might be of finite length but not of finite mean length. However, this is not the case, at least in dimension $d = 2$. We begin to show this by first establishing the finite mean length of constrained Π-geodesics, restricted to lie within specified (two-dimensional) balls.

LEMMA 5.1. *Suppose $d = 2$ and $\gamma > 2$. Consider $x_1, x_2 \in \mathbb{R}^2$, fix $r_0 > |x_2 - x_1|$, and consider the least time by which it is possible to connect x_1 to x_2 by a Π-path which remains entirely within* $\mathrm{ball}(\frac{x_1+x_2}{2}, r_0)$:

$$T^*_{r_0} = \inf\Big\{ T : \text{ there is } \xi \in \mathcal{A}_T \text{ such that } \xi(0) = x_1, \xi(T) = x_2, \text{ and } \xi(t) \in \mathrm{ball}\Big(\frac{x_1 + x_2}{2}, r_0\Big) \text{ for all } t \in [0, T] \Big\}.$$

Then $\mathbb{E}[T^*_{r_0} | V_0] < \infty$, where V_0 is the speed limit of the fastest line hitting $\mathrm{ball}(\frac{x_1+x_2}{2}, r_0)$; *moreover, the following finite expectation provides an upper bound on the mean length of a Π-path connecting x_1 to x_2 within* $\mathrm{ball}(\frac{x_1+x_2}{2}, r_0)$:

$$\mathbb{E}\big[V_0 T^*_{r_0}\big] < \infty.$$

PROOF. Because we work only in dimension 2, and seek an upper bound on Π-geodesic length, we are able to concentrate on Π-paths defined by joining together sequences of line segments; there is no need to negotiate the complexities of the tree construction described in Theorems 3.1 and 3.6. The time taken by such a Π-path, constrained to lie within $\mathrm{ball}(\frac{x_1+x_2}{2}, r_0)$ and connecting x_1 to x_2, necessarily provides an upper bound on the constrained Π-geodesic lying within $\mathrm{ball}(\frac{x_1+x_2}{2}, r_0)$ and connecting x_1 to x_2. Thus the finiteness of $\mathbb{E}[V_0 T^*_{r_0}]$, together with the fact that V_0 is the maximum speed attainable within $\mathrm{ball}(\frac{x_1+x_2}{2}, r_0)$, provides an upper bound on the mean length of the constrained Π-geodesic connecting x_1 to x_2 which is restricted to lie within $\mathrm{ball}(\frac{x_1+x_2}{2}, r_0)$.

Suppose that $|x_2 - x_1| = r_1 < \frac{1}{2}r_0$. Without loss of generality, we suppose that $x_1 + x_2 = \mathbf{o} = (0,0)$, and $x_1 = (-\frac{1}{2}r_1, 0)$, $x_2 = (\frac{1}{2}r_1, 0)$. We shall join x_1 and x_2 together by working toward the two points by two paths commencing on the line segment $\sigma_1 = \{(0,h) : 0 \le h \le r_1\}$; we will then be able to join the two Π-paths together by prolonging the first line segment used in the construction of one of the Π-paths.

The constructions of the two Π-paths are entirely similar, so we shall focus on the Π-path leading to x_1.

Suppose that the fastest line intersecting ball$(\mathbf{o}, r_0)$ has speed limit V_0. Exploiting the notion of meta-slowness described above in the proof of Theorem 2.6, we know that if $S_0 = V_0^{-(\gamma-1)}$ is the meta-slowness of this line then

$$S_0 = \frac{1}{\pi r_0} E_0 \qquad \text{where } E_0 \text{ is distributed as Exponential(1).} \tag{5.1}$$

The following formulae are simplified if our Π-path constructions are required to avoid using this line. The first line used in the Π-path running to x_1 will be the fastest line ℓ_1 with speed less than V_0 and intersecting both the line segment σ_1 and the line segment from x_1 to $\frac{3}{2}x_1$. Suppose that the speed limit of this line is V_1, so the meta-slowness is $S_1 = V_1^{-(\gamma-1)}$. We use standard integral geometry of lines and Pythagoras to show that the line measure of all such lines is

$$\tfrac{1}{2}\big(|x_1 - (0,r_1)| + |\tfrac{3}{2}x_1| - |\tfrac{3}{2}x_1 - (0,r_1)| - |x_1|\big) = \tfrac{\sqrt{5}-2}{4} r_1.$$

Consequently, we may deduce

$$S_1 = S_0 + \frac{4}{(\sqrt{5}-2)r_1} E_1,$$

where E_1 is distributed as Exponential(1) and E_1 is independent of S_0 (equivalently V_0) and the geometry of the line ℓ_1.

Let y_1 be the point on ℓ_1 closest to x_1, and note that the distance from y_1 to σ_1 along ℓ_1 is bounded above by the distance from $\frac{3}{2}x_1$ to $(0, r_1)$, namely

$$\sqrt{\tfrac{9}{16}r_1^2 + r_1^2} = \tfrac{5}{4} r_1.$$

The construction is illustrated in Figure 8.

This construction is continued recursively, for example, replacing the origin by the point y_1 on ℓ_1 closest to x_1, r_0 by r_1, and replacing the segment σ_1 by a segment σ_2 begun at y_1, directed along ℓ_1 toward the start of the Π-path, of length $r_2 = 2|y_1 - x_1|$. Simple geometric arguments show that both $|y_1 - x_1|$ and $|y_1 - \frac{3}{2}x_1|$ are bounded above by $\frac{1}{4}r_1$, so the distance that σ_2 extends from $\frac{3}{2}x_1$ cannot exceed $\frac{3}{4}r_1$, while the distance between $\frac{3}{2}x_1$ and σ_1 is $\frac{3}{4}r_1$. This construction can be used to generate a new line ℓ_2, of meta-slowness $S_2 = V_2^{-(\gamma-1)}$ which is

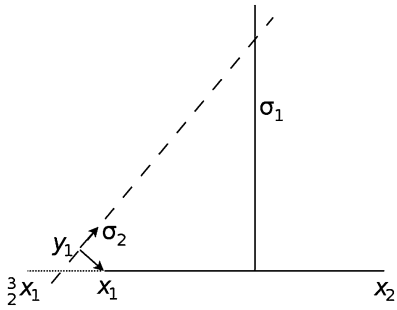

FIG. 8. *Successive construction of one of two components of a* Π*-path from* x_1 *to* x_2.

required to be strictly greater than S_1, and a new closest distance $r_2/2$ from x_1 to ℓ_2. The calculations show that $r_2 \le r_1/2$.

In general, the nth line ℓ_n of the construction has meta-slowness $S_n = V_n^{-(\gamma-1)}$ with

$$S_n = S_{n-1} + \frac{4}{(\sqrt{5}-2)r_n} E_n, \tag{5.2}$$

where E_n is distributed as Exponential(1) and T_1 is independent of $S_0, \dots, S_{n-1}$ (equivalently $V_0, \dots, V_{n-1}$) and the geometry of the lines $\ell_1, \dots, \ell_{n-1}$. Here, r_n is the closest distance from x_1 to ℓ_n, and $r_n < r_{n-1}/2$; the length of the new segment (running from σ_n to y_n along ℓ_n) is bounded above by $\frac{5}{4} r_n$.

Evidently, we have constructed a Π-path from σ_1 to x_1, built as a sequence of line segments. Total time of travel is bounded above by

$$\begin{aligned}
&\sum_{n=1}^{\infty} S_n^{1/(\gamma-1)} \times \frac{5}{4} r_n \\
&\quad = \frac{5}{4} \sum_{n=1}^{\infty} \left(S_0 + \frac{4}{\sqrt{5}-2} \left(\frac{E_1}{r_1} + \cdots + \frac{E_n}{r_n} \right) \right)^{1/(\gamma-1)} r_n \\
&\quad \le \frac{5 r_0^{(\gamma-2)/(\gamma-1)}}{4} \\
&\qquad \times \sum_{n=1}^{\infty} \Bigg(2^{-(n-1)} r_0 S_0 \\
&\qquad\qquad + \frac{4}{\sqrt{5}-2} \left(2^{-(n-1)} E_1 + 2^{-(n-2)} E_2 + \cdots + E_n \right) \Bigg)^{1/(\gamma-1)} \\
&\qquad \times \left(2^{(\gamma-2)/(\gamma-1)} \right)^{-n},
\end{aligned}$$

where the second step uses $r_n < r_{n-1}/2$ and $r_1 < r_0$. We can use the conditional Jensen's inequality for the concave function $u \mapsto u^{1/(\gamma-1)}$ (concave because

$\gamma > 2$) to deduce that the mean total time of travel, conditional on V_0 (equivalently S_0), is bounded above by

$$\frac{5r_0^{(\gamma-2)/(\gamma-1)}}{4}\sum_{n=1}^{\infty}\left(2^{-(n-1)}r_0S_0+\frac{8}{\sqrt{5}-2}\right)^{1/(\gamma-1)}\left(2^{(\gamma-2)/(\gamma-1)}\right)^{-n}. \tag{5.3}$$

Comparison with a geometric sum shows that this sum is finite, since $\gamma > 2$.

We deduce the finiteness of the conditional mean time from x_1 to x_2, since the path can be completed by extending either ℓ_1 or its counterpart in the x_2 path construction; the extra length required is bounded above by $r_1 < r_0$, and the extra time required is therefore bounded above by $V_0^{-1}r_0$.

Finally, finiteness of mean length follows by multiplying the conditional mean time by $V_0 = S_0^{-1/(\gamma-1)}$ and then taking the expectation. The decisive calculation concerns what happens to the conditional bound (5.3) when multiplying through by $S_0^{-1/(\gamma-1)}$ and taking the expectation; we obtain a mean length upper bound of

$$\begin{aligned}
&\mathbb{E}\left[\left(\frac{1}{S_0}\right)^{1/(\gamma-1)}\frac{5r_0^{(\gamma-2)/(\gamma-1)}}{4}\sum_{n=1}^{\infty}\left(2^{-(n-1)}r_0S_0+\frac{8}{\sqrt{5}-2}\right)^{1/(\gamma-1)}\right.\\
&\qquad\left.\times\left(2^{(\gamma-2)/(\gamma-1)}\right)^{-n}\right]\\
&=\frac{5r_0^{(\gamma-2)/(\gamma-1)}}{4}\sum_{n=1}^{\infty}\mathbb{E}\left[\left(\frac{1}{S_0}\right)^{1/(\gamma-1)}\left(2^{-(n-1)}r_0S_0+\frac{8}{\sqrt{5}-2}\right)^{1/(\gamma-1)}\right]\\
&\qquad\times\left(2^{(\gamma-2)/(\gamma-1)}\right)^{-n}\\
&\le\frac{5r_0^{(\gamma-2)/(\gamma-1)}}{4}\sum_{n=1}^{\infty}\mathbb{E}\left[\left(2^{-(n-1)}r_0S_0+\frac{8}{\sqrt{5}-2}\right)^{1/(\gamma-)1};S_0\ge 1\right]\\
&\qquad\times\left(2^{(\gamma-2)/(\gamma-1)}\right)^{-n}\\
&\quad+\frac{5r_0^{(\gamma-2)/(\gamma-1)}}{4}\mathbb{E}\left[\left(\frac{1}{S_0}\right)^{1/(\gamma-1)};S_0<1\right]\\
&\quad\times\sum_{n=1}^{\infty}\left(2^{-(n-1)}r_0+\frac{8}{\sqrt{5}-2}\right)^{1/(\gamma-1)}\left(2^{(\gamma-2)/(\gamma-1)}\right)^{-n}.
\end{aligned} \tag{5.4}$$

Finiteness of the first summand follows from the conditional Jensen's inequality as before (note that $\gamma > 2$). Finiteness of the second summand follows by noting, as $\gamma > 2$,

$$\mathbb{E}\left[\left(\frac{1}{S_0}\right)^{1/(\gamma-1)};S_0<1\right]<\infty. \qquad \square$$

We can now prove the full result: the Π-geodesics between specified points are of finite mean length if $d = 2$.

THEOREM 5.2. *Suppose $d = 2$ and $\gamma > 2$. Consider a Π-geodesic ξ connecting two points x_1 and x_2. The mean length of ξ is finite.*

PROOF. Consider two points x_1, x_2. Without loss of generality, set $r_0 > \frac{3}{\sqrt{2}} r_1 = \frac{3}{\sqrt{2}} |x_2 - x_1|$, $\frac{1}{2}(x_1 + x_2) = \mathbf{o}$, and $x_1, x_2 \in \text{ball}(\mathbf{o}, r_0)$. Note that we can pick r_0 as large as we please. We wish to show that the Π-geodesic ξ from x_1 to x_2 is of finite length.

It is immediate from the Π-geodesic property that the time spent by ξ in $\text{ball}(\mathbf{o}, r_0)$ cannot exceed the time spent traveling from x_1 to x_2 using the path described in Lemma 5.1. Following the arguments of Lemma 5.1, we deduce finiteness of mean for the length of the portion of ξ lying in $\text{ball}(\mathbf{o}, r_0)$.

Let $V_0 = S_0^{-1/(\gamma-1)}$ be the fastest line hitting $\text{ball}(\mathbf{o}, r_0)$. Recall from equation (5.1) of Lemma 5.1 that

$$S_0 = \frac{1}{\pi r_0} E_0 \qquad \text{where } E_0 \text{ is distributed as Exponential}(1).$$

On the other hand, consider the "racetrack" around $\mathbf{o}$ formed by the fastest lines slower than V_0 and connecting the short sides of rectangles of sides r_1 and $3r_1$, placed to surround a central $r_1 \times r_1$ square (see Figure 9). By our choice of r_0, the rectangles are all contained in $\text{ball}(\mathbf{o}, r_0)$. Each of these lines intersects the $3r_1 \times 3r_1$ square in a segment of length at most $\sqrt{10} r_1$. Moreover, the invariant line measure of the set of lines joining the short sides of a rectangle of sides r_1 and

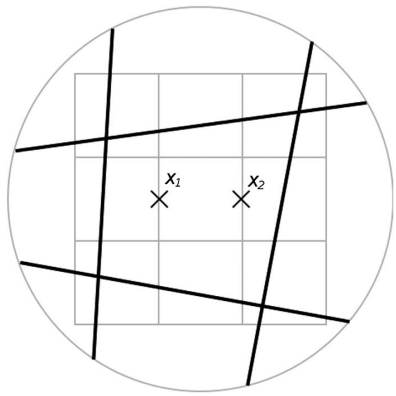

FIG. 9. *Illustration of the racetrack construction: four $r_1 \times 3r_1$ rectangles placed to surround a central $r_1 \times r_1$ square, which is centered inside a disc of radius $r_0 > \frac{3}{\sqrt{2}} r_1$. The racetrack is formed by four lines connecting the short sides of each rectangle, chosen to be the fastest such lines which are strictly slower than the fastest line hitting the disc.*

$3r_1$ is given by

$$\tfrac{1}{2}(2 \times \sqrt{10}r_1 - 2 \times 3r_1) = (\sqrt{10} - 3)r_1.$$

Therefore, the speed limits $V_i' = (S_i')^{-1/(\gamma-1)}$ $(i = 1, 2, 3, 4)$ of these lines have distributions given by

$$S_i' = S_0 + \frac{1}{(\sqrt{10} - 3)r_1} E_i' \qquad \text{where } E_i' \text{ is distributed as Exponential(1).}$$

Here, the E_1', E_2', E_3', E_4' are independent of each other and of S_0; this can be argued based on the facts that they are based on line-sets which are disjoint and conditioned on being slower than S_0. The racetrack establishes a path of length at most $4\sqrt{10}r_1$, which can be traversed in time at most

$$\begin{aligned} T_* &= \sqrt{10}r_1 \sum_{i=1}^{4} \left(S_0 + \frac{1}{(\sqrt{10} - 3)r_1} E_i' \right)^{1/(\gamma-1)} \\ &= \sqrt{10}r_1^{(\gamma-2)/(\gamma-1)} \sum_{i=1}^{4} \left(r_1 S_0 + \frac{1}{\sqrt{10} - 3} E_i' \right)^{1/(\gamma-1)} \\ &\le \sqrt{10}r_1^{(\gamma-2)/(\gamma-1)} \left(4r_1^{1/(\gamma-1)} S_0^{1/(\gamma-1)} + \sum_{i=1}^{4} \left(\frac{1}{\sqrt{10} - 3} E_i' \right)^{1/(\gamma-1)} \right) \\ &= 4\sqrt{10}r_1 S_0^{1/(\gamma-1)} + \sqrt{10}r_1^{(\gamma-2)/(\gamma-1)} \sum_{i=1}^{4} \left(\frac{1}{\sqrt{10} - 3} E_i' \right)^{1/(\gamma-1)}, \end{aligned} \tag{5.5}$$

where the inequality follows from the Minkowski inequality (note that $\gamma > 2$). It follows that ξ cannot spend more than T_* of time outside of ball$(\mathbf{o}, r_0)$, since otherwise it would be possible to take a short-cut involving only some of the racetrack and two portions of ξ lying within ball$(\mathbf{o}, r_0)$, thus traveling from x_1 to x_2 in less time overall.

We now apply the comparison technique used in the proof of Theorem 2.6, using a scalar comparison process y. We suppose that ξ starts at $\partial\,\text{ball}(\mathbf{o}, r_0)$, so $\text{dist}(\xi(0), \mathbf{o}) = r_0$. Then $|\xi| < y$, where $y(0) = 0$ and $y'(t) = \overline{V}(y(t))$. Note that the fastest line hitting ball$(\mathbf{o}, r_0)$ has speed limit V_0, so $\overline{V}(y(0)) = V_0$. Moreover, $\overline{V}(r)$ for $r > r_0$ is based entirely on lines with speeds faster than V_0, and is therefore independent of E_1', E_2', E_3', E_4'.

In dimension $d = 2$, generalized distance is simply ordinary distance. So the recursive formulation (2.6) becomes

$$R_n - R_{n-1} = \frac{1}{S_{n-1}} \text{Exponential}(\pi), \tag{5.6}$$

$$S_n = S_{n-1} U_n, \tag{5.7}$$

for independent Uniform$(0,1)$ random variables U_i, with distribution of S_0 as above.

The times between successive changes of speed are given by

$$S_{n-1}^{1/(\gamma-1)}(R_n - R_{n-1}) = S_{n-1}^{-(\gamma-2)/(\gamma-1)}\,\text{Exponential}(\pi).$$

We know that S_n decreases as $n \to \infty$. Accordingly, a coupling argument shows that the number N_{T_*} of changes of speed by time T_* will not exceed $\widetilde{N}$, where $\widetilde{N}$ has distribution Poisson$(\pi S_0^{(\gamma-2)/(\gamma-1)} T_*)$ when conditioned on S_0 and T_*, and is independent of the actual changes of speed (though not of T^* or S_0). Thus, the final speed is no more than

$$S_0^{-1/(\gamma-1)} \prod_{n=1}^{\widetilde{N}} U_n^{-1/(\gamma-1)},$$

and the distance travelled by the Π-geodesic outside ball$(\mathbf{o}, r_0)$ cannot exceed

$$S_0^{-1/(\gamma-1)} T_* \prod_{n=1}^{\widetilde{N}} U_n^{-1/(\gamma-1)}.$$

Conditioning on E_1', E_2', E_3', E_4' and S_0, we can integrate out first the U_i's and then the Poissonian variation $\widetilde{N}$ from the resulting bound on mean distance traveled, and then use the upper bound on T_* specified by (5.5). We thus obtain the following bound on mean distance traveled, using $S_0 = \frac{1}{\pi r_0} E_0$ and exchangeability of the E_i':

$$\begin{aligned}
&\mathbb{E}\left[S_0^{-1/(\gamma-1)} T_* \times \prod_{n=1}^{\widetilde{N}} U_n^{-1/(\gamma-1)} \right] \\
&\quad = \mathbb{E}\left[S_0^{-1/(\gamma-1)} T_* \times \left(\frac{\gamma-1}{\gamma-2} \right)^{\widetilde{N}} \right] \\
&\quad = \mathbb{E}\left[S_0^{-1/(\gamma-1)} T_* \times \exp\left(\frac{\pi S_0^{(\gamma-2)/(\gamma-1)}}{\gamma-2} T_* \right) \right] \\
&\quad \le 4\sqrt{10}\,\mathbb{E}\left[\left(r_1 + S_0^{-1/(\gamma-1)} r_1^{(\gamma-2)/(\gamma-1)} \left(\frac{1}{\sqrt{10}-3} E_1' \right)^{1/(\gamma-1)} \right) \right. \\
&\qquad\qquad \times \exp\left(\frac{4\sqrt{10}\pi r_1 S_0}{\gamma-2} \right) \\
&\qquad\qquad \left. \times \exp\left(\frac{\sqrt{10}\pi (r_1 S_0)^{(\gamma-2)/(\gamma-1)}}{\gamma-2} \sum_{i=1}^{4} \left(\frac{1}{\sqrt{10}-3} E_i' \right)^{1/(\gamma-1)} \right) \right] \\
&\quad \le 4\sqrt{10} r_1 \mathbb{E}\left[\left(1 + \left(\frac{\pi}{\sqrt{10}-3} \right)^{1/(\gamma-1)} \left(\frac{r_0}{r_1} \right)^{1/(\gamma-1)} \left(\frac{E_1'}{E_0} \right)^{1/(\gamma-1)} \right) \right.
\end{aligned}$$

$$\times \exp\left(\frac{4\sqrt{10}}{\gamma-2}\frac{r_1}{r_0}E_0\right)$$
$$\times \exp\left(\frac{\sqrt{10}}{\gamma-2}\frac{\pi^{1/(\gamma-1)}}{(\sqrt{10}-3)^{1/(\gamma-1)}}\left(\frac{r_1}{r_0}\right)^{(\gamma-2)/(\gamma-1)}\right.$$
$$\left.\times \sum_{i=1}^{4} E_0^{(\gamma-2)/(\gamma-1)}\left(E_i'\right)^{1/(\gamma-1)}\right)\Bigg].$$

But now we can apply the following simple inequality (for $E_0 > 0$, $E_i' > 0$):

$$(E_0)^{(\gamma-2)/(\gamma-1)}\left(E_i'\right)^{1/(\gamma-1)} = (E_0)^{1-1/(\gamma-1)}\left(E_i'\right)^{1/(\gamma-1)} \le E_0 + E_i',$$

and so deduce that

$$\mathbb{E}\left[S_0^{-1/(\gamma-1)}T_* \times \prod_{n=1}^{N_{T_*}} U_n^{-1/(\gamma-1)}\right]$$
$$\le 4\sqrt{10}r_1\mathbb{E}\Bigg[\left(1+\left(\frac{\pi}{\sqrt{10}-3}\right)^{1/(\gamma-1)}\left(\frac{r_0}{r_1}\right)^{1/(\gamma-1)}\left(\frac{E_1'}{E_0}\right)^{1/(\gamma-1)}\right)$$
$$\times \exp\left(\frac{4\sqrt{10}}{\gamma-2}\frac{r_1}{r_0}E_0\right) \tag{5.8}$$
$$\times \exp\left(\frac{\sqrt{10}}{\gamma-2}\frac{\pi^{1/(\gamma-1)}}{(\sqrt{10}-3)^{1/(\gamma-1)}}\left(\frac{r_1}{r_0}\right)^{(\gamma-2)/(\gamma-1)}\right.$$
$$\left.\times \left(4E_0 + E_1' + E_2' + E_3' + E_4'\right)\right)\Bigg].$$

Now the expectation can be bounded above by an expression involving finite Gamma integrals of the forms

$$\int_0^\infty \exp(-\beta u)\,\mathrm{d}u, \qquad \int_0^\infty u^{1/(\gamma-1)}\exp(-\beta u)\,\mathrm{d}u,$$
$$\int_0^\infty u^{-1/(\gamma-1)}\exp(-\beta u)\,\mathrm{d}u,$$

for $\beta > 0$ (once r_0 is chosen sufficiently large) and $\gamma > 2$. Consequently, the mean length of the Π-geodesic outside of ball$(\mathbf{o}, r_0)$ must also be finite, proving the theorem. □

This work shows that planar spatial networks formed from the Poisson line process model satisfy property 3 of Definition 1.1.

6. Further properties of Π-geodesics in $d = 2$. Finally, we show that in dimension $d = 2$ any specified point x almost surely possesses just one Π-geodesic

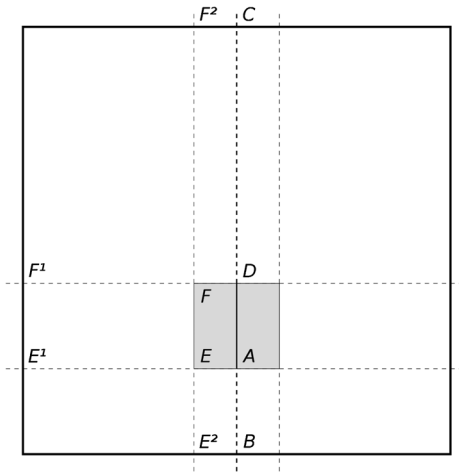

FIG. 10. *Construction forcing certain* Π*-geodesics to pass through the points A and B.*

to ∞; moreover, that for any three distinct points $x, y, z \in \mathbb{R}^2$ almost surely the Π-geodesics from x to y and from x to z coincide for a nontrivial initial segment; and also that if Ξ is an independent Poisson point process in $\mathbb{R}^2$ then almost surely the totality of all Π-geodesics between points of Ξ forms a fiber process [13], Section 8.3, which places finite total length in any given compact subset of $\mathbb{R}^2$.

This last result shows that the network generated by Π possesses a very weak variant of Aldous' *SIRSN property* [2, 3]; a SIRSN (scale-invariant random spatial network) would have the property that the *mean* total length per unit area was finite (weak SIRSN property), and moreover the property that the mean total length of connecting routes of distance at least 1 from start and source would remain bounded as the intensity of Ξ increased to infinity. It is conjectured that the network generated by Π is a true SIRSN; however at present all we can prove is the above "pre-SIRSN" property (but see the remark about Jonas Kahn's observations in Section 7 below).

All three of these results depend on the same construction, based on Aldous [2], Figure 6: consider the behavior of Π-geodesics starting from points in a 2×2 square centered on the origin and ending outside a 10×10 square centered at $(0, 2)$. Condition on the sides of the two squares and the y-axis all being subsets of lines from Π, with speeds as follows: the sides of the 2×2 square have speed a, the y-axis has speed b, and the sides of the 10×10 square have speed c. Suppose further that no other lines of Π_1 (speed exceeding 1) hit the 10×10 square. Figure 10 illustrates the construction.

LEMMA 6.1. *In the above situation, suppose that* $c > 10b > 59a/3 > 354/3$. *Then any* Π*-geodesic connecting the interior of the* 2×2 *square and the exterior of the* 10×10 *square must pass through the points* $A = (0, -1)$ *and* $B = (0, -3)$.

PROOF. To simplify exposition, we can and shall confine our attention to Π-geodesics constrained to lie in or on the 10×10 square.

First, note from the figure that the construction can be divided into rectangles of dimensions 2×2, 1×2, 1×6, 2×4 and 6×4. For each of these rectangles, the sides have speed at least c, while any other lines intersecting the rectangles have speed not exceeding 1. Geometric comparisons show that, for any of these squares, Π-geodesics between pairs of points on the perimeter cannot intersect the interior. This is a "no short-cut" condition for Π-geodesics. In particular, a (constrained) Π-geodesic from the 2×2 square to point C must be confined to the union of the 2×2 square and the other square boundaries.

Consequently, such a Π-geodesic must have a final segment which is one of

1. $A \to B \to C$ (the $B \to C$ part using the perimeter of the 10×10 square);
2. $D \to C$;
3. $E \to E^2 \to C$ (last part using perimeter);
4. $E \to E^1 \to C$ (last part using perimeter);
5. $F \to F^2 \to C$ (last part using perimeter);
6. $F \to F^1 \to C$ (last part using perimeter);
7. or one of four cases which are mirror images of cases 3–6.

We can now compare times taken by these alternative routes: under the condition $c > 10b > 59a/3 > 354/3$ it transpires that the quickest route always passes through the points A and B as required. □

This lemma enables soft proofs of the three theorems of this section.

THEOREM 6.2. *Suppose $\gamma > d = 2$. With probability* 1, *for any point x there is one and only one Π-geodesic from x to ∞; moreover all such infinite Π-geodesics eventually coalesce when sufficiently far away from the origin.*

PROOF. Small perturbations of the structure described in Lemma 6.1 will have the same property (Π-geodesics from within small squares to exteriors of large squares all pass through specified points), and so there is a positive probability $\varepsilon > 0$ that Π will generate a structure ensuring that all Π geodesics from the 2×2 square reaching out further than the 10×10 square will have to pass through a specified pair of points near to A and B.

Moreover, we can use scale-invariance to generate further structures at larger scales, such that whether or not a corresponding perturbation of each structure is realized is independent of whether or not the other structures are realized. An appeal to the second Borel–Cantelli lemma then shows that there must be an infinite sequence of planar points $(0, -3a_1)$, $(0, -3a_2), \ldots \to \infty$, such that if $x \in [-a_n, a_n]^2$ and $y \notin [-5a_n, 5a_n] \times [-3a_n, 7a_n]$ then any Π-geodesic from y to x must pass through $(0, -3a_n)$. Moreover, the section of this Π-geodesic from x to $(0, -3a_n)$ is uniquely determined, since for almost all y the Π-geodesic from

y to x will be unique (Theorem 4.4). Successive sections of similar Π-geodesics therefore build up a unique Π-geodesic from x to ∞. Moreover, the nature of the structure described in Lemma 6.1 ensures that, for any other planar point y, this other point will be included in the smaller of the two rectangles of structures at sufficiently large scales: eventual coalescence of all infinite Π-geodesics thus follows. □

We can now establish a result similar to that of [10] for the planar Brownian map: almost surely Π-geodesics emanating from a given point must initially coalesce. (Evidently, this cannot hold for *all* points: consider points actually lying on a Π-geodesic.)

THEOREM 6.3. *Suppose $\gamma > d = 2$. Almost surely for any distinct points x, y, z, the Π-geodesics from x to y and from x to z coincide for a nontrivial initial segment.*

PROOF. The argument follows that of Theorem 6.2, except that structures are now generated at increasingly smaller scales, all surrounding x. □

THEOREM 6.4. *Suppose $\gamma > d = 2$. The network generated by Π has the* pre-SIRSN property, *in the sense that if Ξ is an independent Poisson point process in $\mathbb{R}^2$ then almost surely the totality of all Π-geodesics between points of Ξ intersected with a compact set has finite total length.*

PROOF. By scaling and monotonicity, we may suppose that the compact set in question is the 1×1 square centred at the origin. Arguing as in Theorem 6.3, at a suitably large scale there will be a structure which forces all Π-geodesics between points in the (1×1) square and points in the exterior of a $L \times L$ square to pass through a specified point H near the boundary of the $L \times L$ square. Here, L is random but depends only on Π, not Ξ.

Let N be the random number of points placed by Ξ in the $L \times L$ square. Then at most $\binom{N+1}{2}$ Π-geodesics can intersect the 1×1 square (based on these points and on H). Each of these Π-geodesics has finite length (Theorem 5.2), so the result follows. □

This proves property 4 of Definition 1.1 for the planar case.

7. Conclusion. This paper has established:

1. Basic metric space properties of the Poisson line process model, including existence of minimum-time paths;
2. Extension of the metric space properties of the Poisson line process model to higher dimensions;

3. Approximation results for minimum-time paths ("Π-geodesics");
4. Almost-sure uniqueness and finite mean length of Π-geodesics in the planar case;
5. Local finiteness of resulting networks in the planar case.

As a result, it follows that the planar Poisson line process model produces a pre-SIRSN (Definition 1.1). The major outstanding question is, whether in fact the weak SIRSN or even full SIRSN properties hold for the planar Poisson line process model. Extending the method of Theorem 6.4 would require a much more quantitative approach; it would be necessary to estimate the scale at which there would exist structures forcing large Π-geodesics to pass through specified points. Remarkably, Jonas Kahn has recently shown me an argument which establishes the full SIRSN property for the case $d = 2$, using a careful development of the proofs of Theorems 3.1 and 3.6; moreover, his arguments establish property 1 (almost-sure uniqueness of Π-geodesics) and property 6 for *all* dimensions $d \geq 2$.

A linked question concerns the nature of Π-geodesics in the planar case: can they be represented using sequences of line segments from the Poisson line process, or do they necessarily involve the tree-like representations described in the proofs of Theorems 3.1 and 3.6? The methods of Section 4 are suggestive that the answer is yes, but do not entirely exclude the possibility of slow-down at points not lying on Π. One must show that Π-geodesics between pairs of points can *never* slow down to zero speed *en route*. This is conceptually linked to the notion of network transit points [8], as discussed in Aldous [2].

The paper has focused throughout on results relating directly to the pre-SIRSN property of the Poisson line process model. We have noted in passing and without proof some computations which establish sharpness of conditions on γ in our results (Remarks 2.7, 3.2); similar calculations show that the topology of Euclidean space $\mathbb{R}^d$ viewed as a Π-geodesic metric space (for $\gamma > d$) is the normal Euclidean topology.

Finally, we note that an alternative motivation for the above work is given by recent developments in the study of Brownian maps; the random metric space given here can be compared to the Brownian map (note, e.g., that both situations exhibit coalescence of geodesics) and promises by its constructive nature to be (relatively) more amenable to rigorous mathematical investigation, as well as providing higher dimensional constructions. It would be of great interest to clarify the extent to which the two theories can be linked. Note in particular the intriguing prospect of mimicking the Brownian map theory by the construction of "Liouville Brownian motions," perhaps using Dirichlet form theory (compare Berestycki [9], Garban, Rhodes and Vargas [15]).

Acknowledgements. My thanks to David Aldous, who challenged me to try to understand the Poisson line process model for a scale-invariant network.

REFERENCES

[1] ABBASPOUR, H. and MOSKOWITZ, M. (2007). *Basic Lie Theory*. World Scientific, Hackensack, NJ. MR2364699

[2] ALDOUS, D. (2014). Scale-invariant random spatial networks. *Electron. J. Probab.* **19** 41. MR3164768

[3] ALDOUS, D. and GANESAN, K. (2013). True scale-invariant random spatial networks. *Proc. Natl. Acad. Sci. USA* **110** 8782–8785. MR3082274

[4] ALDOUS, D. J. and BARLOW, M. T. (1981). On countable dense random sets. In *Seminar on Probability, XV* (*Univ. Strasbourg, Strasbourg*, 1979/1980) (*French*). *Lecture Notes in Math.* **850** 311–327. Springer, Berlin. MR0622573

[5] ALDOUS, D. J. and KENDALL, W. S. (2008). Short-length routes in low-cost networks via Poisson line patterns. *Adv. in Appl. Probab.* **40** 1–21. MR2411811

[6] AMBARTZUMIAN, R. V. (1990). *Factorization Calculus and Geometric Probability*. *Encyclopedia of Mathematics and Its Applications* **33**. Cambridge Univ. Press, Cambridge. MR1075011

[7] BACCELLI, F., TCHOUMATCHENKO, K. and ZUYEV, S. (2000). Markov paths on the Poisson–Delaunay graph with applications to routing in mobile networks. *Adv. in Appl. Probab.* **32** 1–18. MR1765174

[8] BAST, H., FUNKE, S., SANDERS, P. and SCHULTES, D. (2007). Fast routing in road networks with transit nodes. *Science* **316** 566. MR2313317

[9] BERESTYCKI, N. (2013). Diffusion in planar Liouville quantum gravity. Available at arXiv:1301.3356.

[10] BETTINELLI, J. (2014). Geodesics in Brownian surfaces (Brownian maps). Available at arXiv:1401.3602.

[11] BROUTIN, N., DEVILLERS, O. and HEMSLEY, R. (2014). Efficiently navigating a random Delaunay triangulation. Available at arXiv:1402.6148.

[12] BURAGO, D., BURAGO, Y. and IVANOV, S. (2001). *A Course in Metric Geometry*. *Graduate Studies in Mathematics* **33**. Amer. Math. Soc., Providence, RI. MR1835418

[13] CHIU, S. N., STOYAN, D., KENDALL, W. S. and MECKE, J. (2013). *Stochastic Geometry and Its Applications*. Wiley, Chichester.

[14] EVANS, L. C. (1998). *Partial Differential Equations*. *Graduate Studies in Mathematics* **19**. Amer. Math. Soc., Providence, RI. MR1625845

[15] GARBAN, C., RHODES, R. and VARGAS, V. (2013). Liouville Brownian motion. Available at arXiv:1301.2876.

[16] HAMBLY, B. and LYONS, T. (2010). Uniqueness for the signature of a path of bounded variation and the reduced path group. *Ann. of Math.* (2) **171** 109–167. MR2630037

[17] KENDALL, W. S. (2014). Return to the Poissonian city. *J. Appl. Probab.* **51A** 297–309.

[18] KENDALL, W. S. (2000). Stationary countable dense random sets. *Adv. in Appl. Probab.* **32** 86–100. MR1765169

[19] KENDALL, W. S. (2011). Geodesics and flows in a Poissonian city. *Ann. Appl. Probab.* **21** 801–842. MR2830605

[20] LAGATTA, T. (2010). Geodesics of random Riemannian metrics. Ph.D. thesis, Univ. Arizona, Tuscon, AZ. MR2941350

[21] LOOMIS, L. H. (1953). *An Introduction to Abstract Harmonic Analysis*. Van Nostrand, Toronto. MR0054173

[22] MEYN, S. P. and TWEEDIE, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer, London. MR1287609

[23] SANTALÓ, L. A. (1976). *Integral Geometry and Geometric Probability*. Addison-Wesley, Reading, MA. MR0433364

[24] SWIFT, J. (1733). The Poems of Jonathan Swift, D.D. Project Gutenberg.

[25] VERVAAT, W. (1979). On a stochastic difference equation and a representation of nonnegative infinitely divisible random variables. *Adv. in Appl. Probab.* **11** 750–783. MR0544194

DEPARTMENT OF STATISTICS
UNIVERSITY OF WARWICK
COVENTRY CV4 7AL
UNITED KINGDOM
E-MAIL: w.s.kendall@warwick.ac.uk
URL: http://www.warwick.ac.uk/wsk