Models for Connected Networks over Random Points and a Route-Length Statistic

David J. Aldous* Julian Shun

September 30, 2009

Abstract

We review mathematically tractable models for connected networks on random points in the plane, emphasising the little-studied class of *proximity graphs* and introducing a new model called the *Hammersley network*. We introduce and motivate a particular statistic R measuring shortness of routes in a network. We show (via Monte Carlo, in part) the trade-off between normalized network length and R in a one-parameter family of proximity graphs. How close this family comes to the optimal trade-off over all possible networks remains an intriguing open question.

`This material has been presented in talks since 2007. It will be periodically updated as technical papers [3, 4, 7, 8] are completed, and ultimately published as some kind of survey.`

*Department of Statistics, 367 Evans Hall # 3860, U.C. Berkeley CA 94720; aldous@stat.berkeley.edu; www.stat.berkeley.edu/users/aldous. Aldous's research supported by N.S.F Grant DMS-0704159.

1 Introduction

This is intended as the least technical in a series of papers [6, 8, 3, 4] on spatial networks. For concreteness, visualize a road network linking a given set of cities. Our first purpose is to introduce a specific statistic R for measuring the effectiveness of a network in providing short routes ("route-length efficiency"). The definition of R and its motivation are given in section 3.2. A second purpose, perhaps of broader interest (hence presented first, in section 2), is to review models for connected networks on deterministic or random points in the plane. Recall that the most studied network model, the *random geometric graph* [33], does not permit both connectivity and bounded normalized length in the $n \to \infty$ limit.. An attractive alternative is the class of *proximity graphs*, reviewed in section 2.3, which in the deterministic case have been studied within computational geometry. Proximity graphs on random points have scarcely been studied, but are potentially interesting for many purposes other than the specific "short route lengths" topic of this paper (see section 6.5), so we take this opportunity to draw the attention of the applied probability community to this class of *random proximity graphs*. One could also imagine constructions which depend on points having specifically the Poisson point process distribution, and one novel such network, which we name the *Hammersley network*, is described in section 2.5.

The central theme in the series of papers is seeking to quantify the tradeoff between network length (precisely, the normalized length L defined at (2)) and route length efficiency statistics such as R. Our particular statistic R is not amenable to explicit calculation even in the "tractable" models of section 2, but in section 4 we present the results from Monte Carlo simulations of these models in the random case. In particular, Figure 7 shows the tradeoff for the particular *β-skeleton* family of proximity graphs.

Given a normalized network length L, for any realization of cities there is some network of normalized length L which minimizes R. As indicated in section 5, by general abstract mathematical arguments there must exist a deterministic function $R_{\mathrm{opt}}(L)$ giving (in the "number of cities $\to \infty$" limit under the random model) the minimum value of R over all possible networks of normalized length L. An intriguing open question is

> how close are the values $R_{\beta\text{-skel}}(L)$ from the β-skeleton proximity graphs to the optimum values $R_{\mathrm{opt}}(L)$?

As discussed in section 5.3, at first sight it looks easy to design heuristic algorithms for networks which should improve over the β-skeletons, e.g. by introducing Steiner points, but in practice we have not succeeded in doing so.

This paper focusses on the random model for city positions because it seems the natural setting for theoretical study. As a complement, in [10] we give empirical data for the values of (L, R) for certain real-world networks (on the 20 largest cities, in each of 10 U.S. States). In [8] we give analytic results and bounds on the trade-off between L and a mathematically more tractable *stretch* statistic $R_{\max}$ at (4), in both worst-case and random-case settings for city positions. Let us also point out a (perhaps) non-obvious insight discussed in section 3.3: in designing networks to be efficient in the sense of providing short routes, the main difficulty is proving short routes between city-pairs at a specific distance (2 - 3 standardized units) apart rather than pairs at large distances apart.

Finally, recall this is a non-technical paper. Our purpose is to elaborate the ideas outlined above; technical aspects may be briefly mentioned (e.g. section 6.2) but will be pursued elsewhere.

2 Models for connected spatial networks

There are several conceptually different ways of defining networks on random points in the plane. To be concrete we call the points *cities*; to be consistent about language we regard x_i as the *position* of city i and represent network edges as line segments (x_i, x_j).

First (sections 2.1 - 2.3) are schemes which use deterministic rules to define edges for an arbitrary deterministic configuration of cities; then one just applies these rules to a random configuration. Second, one can have random rules for edges in a deterministic configuration (e.g. the probability of an edge between cities i and j is a function of Euclidean distance $d(x_i, x_j)$, as in popular *small worlds* models [32]), and again apply to a random configuration. Third, and more subtly, one can have constructions that depend on the randomness model for city positions – section 2.5 provides a novel example.

2.1 The geometric graph

In sections 2.1 - 2.3 we have an arbitrary configuration $\mathbf{x} = \{x_i\}$ of city positions, and a deterministic rule for defining the edge-set $\mathcal{E}$. Usually in graph theory one imagines a *finite* configuration, but note that everything makes sense for *locally finite* configurations too. Where helpful we assume "general position", so that intercity distances $d(x_i, x_j)$ are all distinct.

For the *geometric graph* one fixes $0 < c < \infty$ and defines

$$(x_i, x_j) \in \mathcal{E} \text{ iff } d(x_i, x_j) \leq c.$$

For the K-*neighbor graph* one fixed $K \geq 1$ and defines

> $(x_i, x_j) \in \mathcal{E}$ iff x_i is one of the K closest neighbors of x_j or x_j is one of the K closest neighbors of x_i.

A moment's thought shows these graphs are in general not connected, so we turn to models which are "by construction" connected.

2.2 A nested sequence of connected graphs

For more about the material here and in the next section see [25], which develops algorithmic applications in computational geometry and pattern recognition. But everything we need is immediate from the (careful choice of) definitions. On our arbitrary configuration $\mathbf{x}$ we can define four graphs whose edge-sets are nested as follows:

$$\text{MST} \subseteq \text{relative n'hood} \subseteq \text{Gabriel} \subseteq \text{Delaunay} . \qquad (1)$$

Here are the definitions (for MST and Delaunay, it is easy to check these are equivalent to more familiar definitions). In each case, we write the criterion for an edge (x_i, x_j) to be present.

- *Minimum spanning tree (MST) [22].* There does not exist a sequence $i = k_0, k_1, \ldots, k_m = j$ of cities such that $\max(d(x_{k_0}, x_{k_1}), d(x_{k_1}, x_{k_2}), \ldots, d(x_{k_{m-1}}, x_{k_m})) < d(x_i, x_j)$.
- *Relative neighborhood graph.* There does not exist a city k such that $\max(d(x_i, x_k), d(x_k, x_j)) < d(x_i, x_j)$.
- *Gabriel graph.* There does not exist a city inside the disc with diameter $d(x_i, x_j)$.
- *Delaunay triangulation [21].* There exists some disc, with x_i and x_j on its boundary, so that no city is inside the disc.

The inclusions (1) are immediate from the definitions. Because the MST (for a finite configuration) is connected, all these graphs are connected.

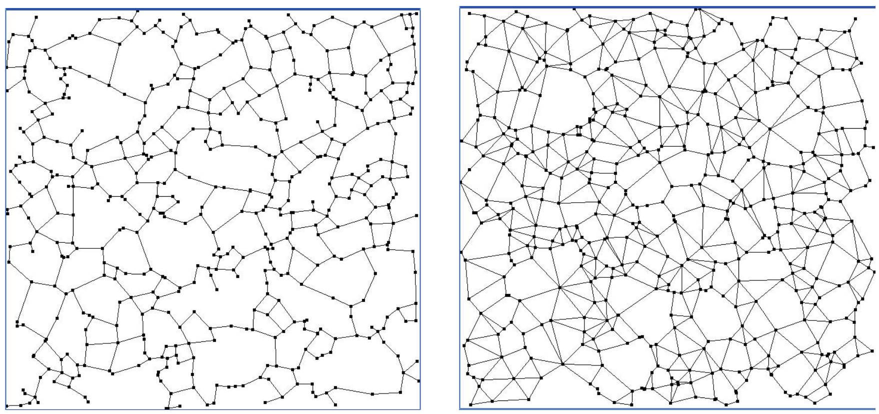

Figure 1. The relative neighborhood graph (left) and Gabriel graph (right) on different realizations of 500 random points.

Figure 1 illustrates the relative neighborhood and Gabriel graphs. Figures for the MST and the Delaunay triangulation can be found online at `http://www.spss.com/research/wilkinson/Applets/edges.html`.

Constructions such as the relative neighborhood and Gabriel graphs have become known loosely as *proximity graphs* in [25] and subsequent literature, and we next take the opportunity to make an implicit definition in the literature into an explicit definition.

2.3 Proximity graphs

Write v_- and v_+ for the points $(-\frac{1}{2}, 0)$ and $(\frac{1}{2}, 0)$. The *lune* is the intersection of the open discs of radii 1 centered at v_- and v_+. So v_- and v_+ are not in the lune but are on its boundary. Define a *template* A to be a subset of $\mathbb{R}^2$ such that

(i) A is a subset of the lune;

(ii) A contains the open line segment (v_-, v_+);

(iii) A is invariant under the "reflection in the y-axis" map $\text{Reflect}_x(x_1, x_2) = (-x_1, x_2)$ and the "reflection in the x-axis" map $\text{Reflect}_y(x_1, x_2) = (x_1, -x_2)$.

(iv) A is open.

For arbitrary points x, y in $\mathbb{R}^2$, define $A(x, y)$ to be the image of A under the transformation (translation, rotation and scaling) that takes (v_-, v_+) to (x, y).

Definition. Given a template A and a locally finite set $\mathcal{V}$ of vertices, the associated *proximity graph* G has edges defined by: for each $x, y \in \mathcal{V}$,

$$(x, y) \text{ is an edge of } G \text{ iff } A(x, y) \text{ contains no vertex of } \mathcal{V}.$$

From the definitions:

- if A is the lune then G is the relative neighborhood graph;
- if A is the disc centered at the origin with radius $1/2$ then G is the Gabriel graph.

But the MST and Delaunay triangulation are *not* instances of proximity graphs.

Note that replacing A by a subset A' can only introduce extra edges. It follows from (1) that the proximity graph is always connected. On the other hand, if A is not a superset of the disc centered at the origin with radius $1/2$, then G might not be a subgraph of the Delaunay triangulation, and in this case edges may cross (i.e. G is not planar).

For a given configuration $\mathbf{x}$, there is a collection of proximity graphs indexed by the template A, so by choosing a monotone one-parameter family of templates one gets a monotone one-parameter family of graphs, analogous to the one-parameter family $\mathcal{G}_c$ of geometric graphs. Here is a popular choice [27] in which $\beta = 1$ gives the Gabriel graph and $\beta = 2$ gives the relative neighborhood graph.

Definition: the β-*skeleton* family.
(i) For $0 < \beta < 1$ let A_β be the intersection of the two open discs of radius $(2\beta)^{-1}$ passing through v_- and v_+.
(ii) For $1 \le \beta \le 2$ let A_β be the intersection of the two open discs of radius $\beta/2$ centered at $(\pm(\beta-1)/2, 0)$.

2.4 Networks based on powers of edge-lengths

It is not hard to think of other ways to define one-parameter families of networks. Here is one scheme, used in e.g. [36]. Fix $1 \le p < \infty$. Given a configuration $\mathbf{x}$, and a route (sequence of vertices) $x_0, x_1, \ldots, x_k$, say the cost of the route is the sum of p'th powers of the step lengths. Now say that a pair (x, y) is an edge of the network $\mathcal{G}_p$ if the cheapest route from x to y is the one-step route. As p increases from 1 to ∞ these networks decrease from the complete graph to the MST. Moreover for $p \ge 2$ the network $\mathcal{G}_p$ is a subgraph of the Gabriel graph.

2.5 The Hammersley network

There is a quite separate recent literature in theoretical probability [23, 24] defining structures such as trees and matchings directly on the infinite Poisson point process. In this spirit, we observe that the *Hammersley process* studied in [5] can be used to define a new network on the infinite Poisson point process, which we name the *Hammersley network*. This network is designed to have the feature that each vertex has exactly 4 edges, in directions NE (between North and East), NW, SE and SW. The conceptual difference from the networks in the previous section is that there is not such a simple "local" criterion for whether a potential edge (x_i, x_j) is in the network. And edges cross, creating junctions.

For a picturesque description, imagine one-eyed frogs sitting on an infinitely long, thin log, each being able to see only the part of the log to their left before the next frog. At random times and positions (precisely, as a space-time Poisson point process of rate 1) a fly lands on the log, at which instant the (unique) frog which can see it jumps left to the fly's position and eats it. This defines a continuous time Markov process (Hammersley process) whose states are the configurations of positions of all the frogs.

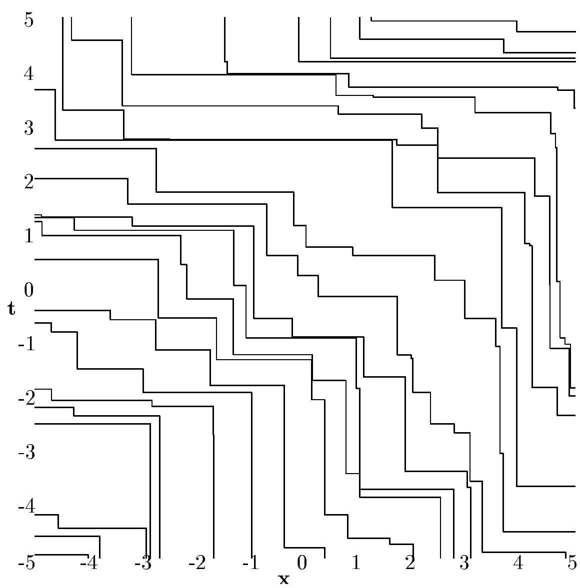

Figure 2. Space-time trajectories in Hammersley's process.

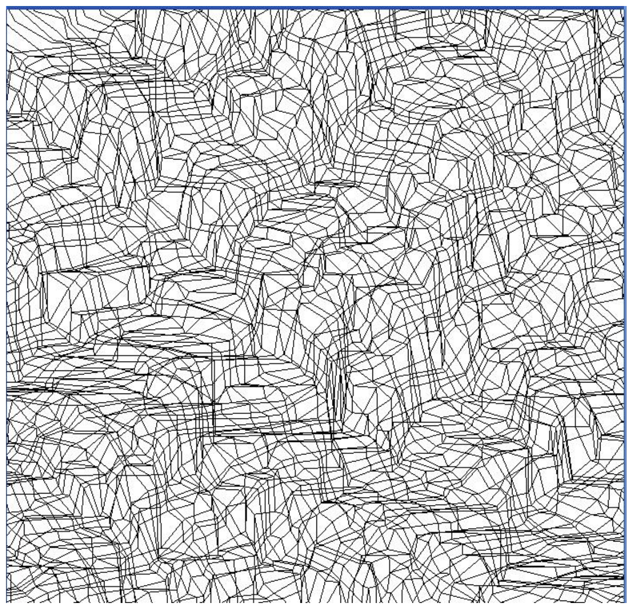

Figure 3. The Hammersley network on 2,500 random points.

There is a *stationary* version of the process in which at each time, the positions of the frogs form a Poisson (rate 1) point process on the line.

Now consider the space-time trajectories of all the frogs, drawn with time increasing upwards on the page. See Figure 2. For each frog, the part of the trajectory between the ends of two successive jumps consists of an upward edge (the frog remains in place as time increases) followed by a leftward edge (the frog jumps left). Reinterpreting the time axis as a second space axis, and introducing compass directions, that part of the trajectory becomes a North edge followed by a West edge. Now replace these two edges by a single North-West straight edge. Doing this procedure for each frog and each pair of successive jumps, we obtain a collection of NW paths; that is, a network in which each city (the reinterpreted space-time random points) has a edge to the NW and an edge to the SE. Finally we repeat with the same realization of the space-time Poisson point process but with frogs jumping rightwards instead of leftwards. This yields a network on the infinite Poisson point process, which we name the *Hammersley network*. See Figure 3.

Remarks. **(a).** To draw the Hammersley network on random points in a *finite square*, one needs external randomization to give the initial (time

0) frog positions, in fact two independent randomizations for the leftwards and the rightwards processes. So to be pedantic, one gets a *random* network over the given realization of cities. However, one can deduce from the theoretical results in [5] that the external randomization has effect only near the boundary of the square.

(b). The property that each vertex has exactly 4 edges, in directions NE (between North and East), NW, SE and SW, is immediate from the construction. Note however that while adjacent NW space-time trajectories in Figure 2 do not cross, the corresponding diagonal roads in the Hammersley network may cross.

2.6 Normalized length

The notion of *normalized network length* L is most easily visualized in the setting of an infinite deterministic network which is "regular" in the sense of consisting of a repeated pattern. First choose the unit of length so that cities have an average density of one per unit area. Then define

$$L = \text{average network length per unit area} \tag{2}$$

$$\bar{d} = \text{average degree (number of incident edges) of cities.} \tag{3}$$

Figure 4 shows the values of L and $\bar{d}$ for some simple "repeated pattern" networks. Though not directly relevant to our study of the random model, we find Figure 4 helpful for two reasons: as intuition for the interpretation of the different numerical values of L, and because we can make very loose analogies (section 6.6) between particular networks on random points and particular deterministic networks.

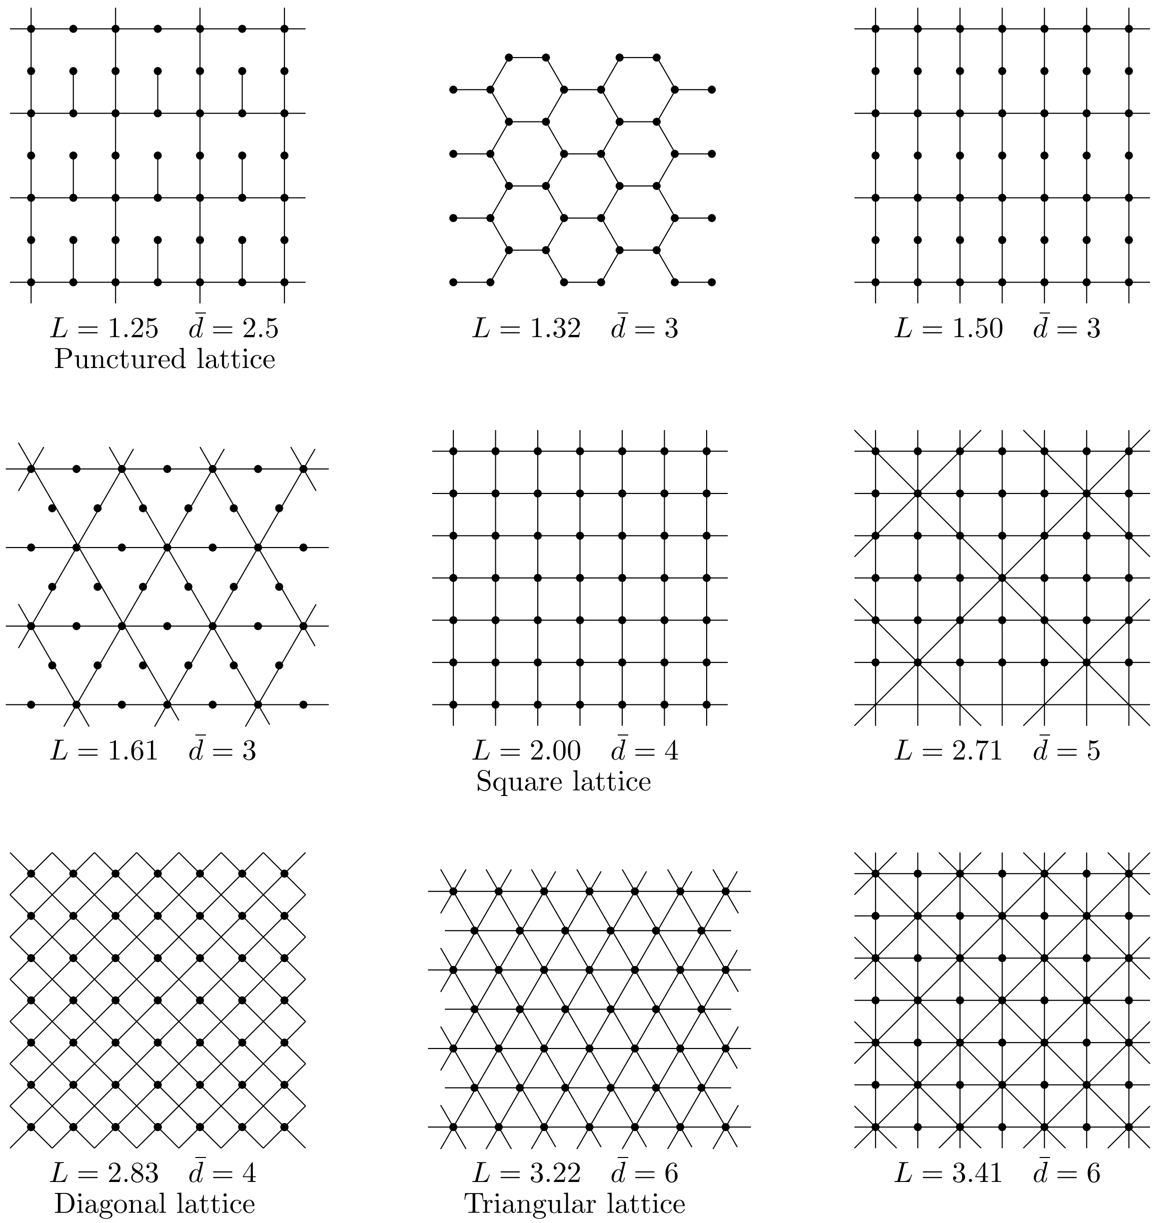

Figure 4. Variant square, triangular and hexagonal lattices. Drawn so that the density of cities is the same in each diagram, and ordered by value of L.

3 Normalized length and route-length efficiency

3.1 The random model

For the remainder of the paper we work with "the random model" for city positions. The *finite model* assumes n random vertices (cities) distributed independently and uniformly in a square of area n. The *infinite model* assumes the Poisson point process of rate 1 (per unit area) in the plane. The quantities $L, \bar{d}$ above and R below that we discuss may be interpreted as exact values in the infinite model or as $n \to \infty$ limits in the finite model: see section 5. We use the word *normalized* as a reminder of the "density 1" convention – we choose the normalized unit of distance to make cities have average density 1 per unit area.

3.2 The route-length efficiency statistic R

In designing a network, it is natural to regard total length as a "cost". The corresponding "benefit" we seek is to have short routes between cities. Write $\ell(i,j)$ for the route-length (length of shortest path) between cities i and j in a given network, and $d(i,j)$ for Euclidean distance between the cities. So $\ell(i,j) \geq d(i,j)$, and we write

$$r(i,j) = \tfrac{\ell(i,j)}{d(i,j)} - 1$$

so that "$r(i,j) = 0.2$" means that route-length is 20% longer than straight line distance. With n cities we get $\binom{n}{2}$ such numbers $r(i,j)$; what is a reasonable way to combine these into a single statistic? Two natural possibilities are

$$\begin{aligned} R_{\text{max}} &:= \max_{j \neq i} r(i,j) \\ R_{\text{ave}} &:= \text{ave}_{(i,j)} r(i,j) \end{aligned} \tag{4}$$

where $\text{ave}_{(i,j)}$ denotes average over all distinct pairs (i,j). The statistic R_{max} has been studied in the context of the design of geometric spanner networks [31] where it is often called the *stretch*. However, being an "extremal" statistic R_{max} seems unsatisfactory as a descriptor of real world networks – for instance, it seems unreasonable to characterize the U.K. rail network as inefficient simply because there is no very direct route between Oxford and Cambridge.

The statistic R_{ave} has a more subtle drawback. Consider a network consisting of

- the minimum-length connected network (Steiner tree) on given cities;
- and a superimposed sparse collection of randomly oriented lines (a *Poisson line process* [38]).

See Figure 5. By choosing the density of lines to be sufficiently low, one can make the normalized network length be arbitrarily close to the minimum needed for connectivity. But it is easy to show (see [6] for careful analysis and stronger result) that one can construct such networks so that $R_{\mathrm{ave}} \to 0$ as $n \to \infty$. Of course no-one would build a road network looking like Figure 5 to link cities, because there are many pairs of nearby cities with only very indirect routes between them. The disadvantage of R_{ave} as a descriptive statistic is that (for large n) most city-pairs are far apart, so the fact that a given network has a small value of R_{ave} says nothing about route-lengths between nearby cities.

We propose a statistic R which is intermediate between R_{ave} and R_{max}. First consider

$$\rho(d) := \text{ mean value of } r(i,j) \text{ over city-pairs with } d(i,j) = d$$

and then define

$$R := \max_{0 \le d < \infty} \rho(d). \tag{5}$$

In words, $R = 0.2$ means that on every scale of distance, route-lengths are on average at most 20% longer than straight line distance.

On an intuitive level, R provides a sensible and intepretable way to compare efficiency of different networks in providing short routes. On a technical level, we see two advantages and one disadvantage of using R instead of R_{ave}.

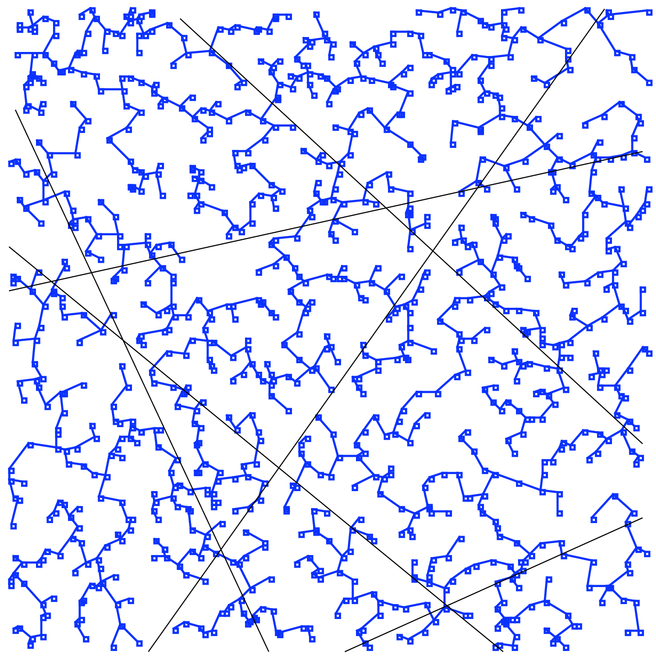

Figure 5. Efficient or inefficient? R_{ave} would judge this network efficient in the $n \to \infty$ limit. [1]

Advantage 1. Using R to measure efficiency, there is a meaningful $n \to \infty$ limit for the network length/efficiency tradeoff (the function $R_{\text{opt}}(L)$ discussed in section 5), and so in particular it makes sense to compare the values of R for networks with different n.

Advantage 2. A more realistic model for traffic would posit that volume of traffic between two cities varies as a power-law $d^{-\gamma}$ of distance d, so that in calculating R_{ave} it would be more realistic to weight by $d^{-\gamma}$. This means that the optimal network, when using R_{ave} as optimality criterion, would depend on γ. Use of R finesses this issue; the value of γ does not affect R.[2]

[1] The tree is actually a MST, taken from `http://www.spss.com/research/wilkinson`, because we could not print a suitable Steiner tree picture, but one can be viewed at `http://www.css.taylor.edu/~bbell/steiner/1600.gif`. Similarly, Table 1 quotes the MST rather than the Steiner tree because we could not find data for normalized length for the Steiner tree.

[2] A related issue is that volume of traffic between two cities should depend on their populations. Intuitively, incorporating random population sizes should make the optimal R *smaller* because the network designer can create shorter routes between larger cities. We see this effect in data [10]; R calculated via population-wieghting is typically smaller.

Disadvantage. The statistic R is tailored to the infinite model, in which it makes sense to consider two cities at exactly distance d apart (then the other city positions form a Poisson point process). For finite n we need to discretize. For the empirical data in [10], where $n = 20$, we average over intervals of width 1 unit (recall the unit of distance is taken such that the density of cities is 1 per unit area); that is for $d = 1, 2, \ldots, 5$ we calculate

$$\tilde{\rho}(d) := \text{ mean value of } r(i,j) \text{ over city-pairs with } d - \tfrac{1}{2} < d(i,j) < d + \tfrac{1}{2}$$

$$\tilde{R} := \max_{1 \le d < \infty} \tilde{\rho}(d) \tag{6}$$

and use $\tilde{R}$ as proxy for R. For larger n we can use shorter intervals. Thus there is in principle a certain fuzziness to the notion of R for finite networks, and in particular it is not clear how to assign a value of R to regular networks such as those in Figure 4. But in practice, for networks we have studied on real-world data and on random points, this is not a problem, as explained next.

3.3 Characteristic shape of the function $\rho(d)$

For the tractable networks on random points (excluding the Hammersley network) the function $\rho(d)$ has a characteristic shape (see Figure 6) attaining its maximum between 2 and 3 and slowly decreasing thereafter. This "smoothness near the maximum" is technically convenient, implying that any calculated value $\tilde{R}$ at (6) is quite insensitive to the choice of discretization.

But we have not tried theoretical study.

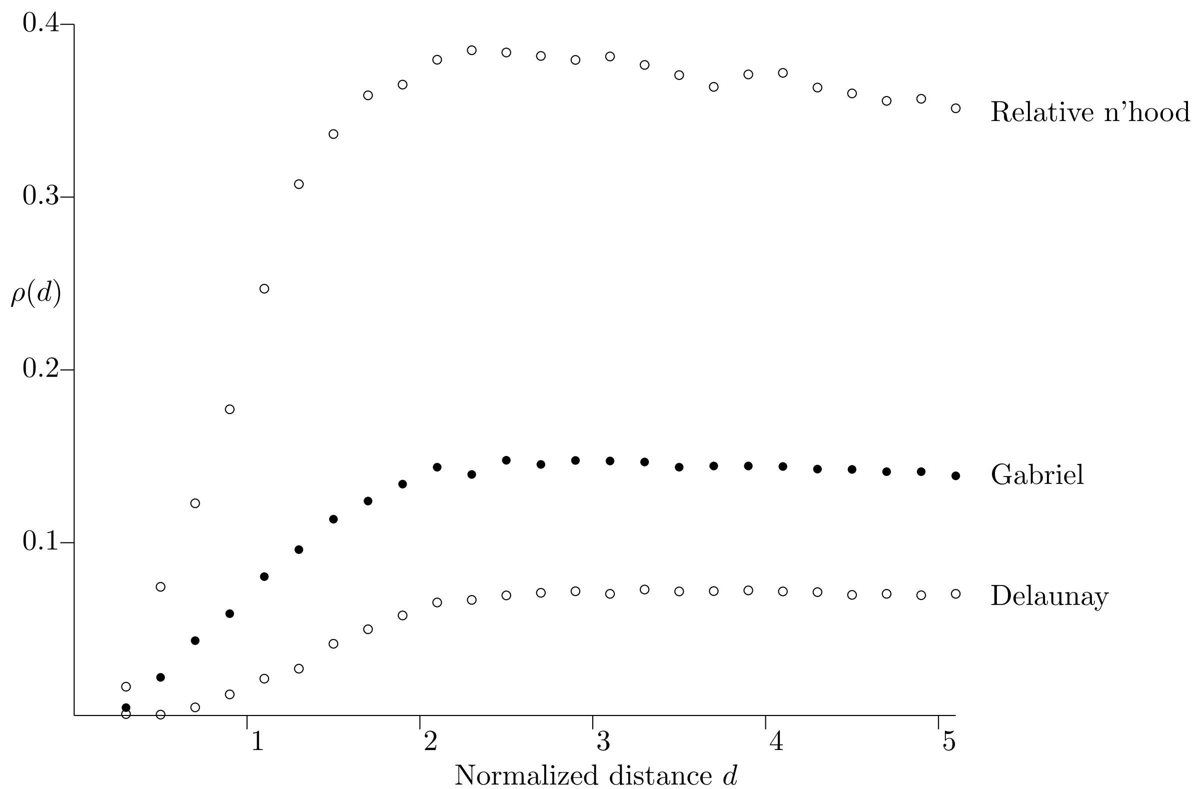

Figure 6. The function $\rho(d)$ for three theoretical networks on random cities. Irregularities are Monte Carlo random variation.

This characteristic shape has a common-sense interpretation. Any efficient network will tend to place roads directly between unusually close city-pairs, implying that $\rho(d)$ should be small for $d < 1$. For large d the presence of multiple alternate routes helps prevent $\rho(d)$ from growing.. At distance $2-3$ from a typical city i there are about $\pi 3^2 - \pi 2^2 \approx 16$ other cities j. For some of these j there will be cities k near the straight line from i to j, so the network designer can create roads from i to k to j. The difficulty arises where there is no such intermediate city k: including a direct road (x_i, x_j) will increase L, but not including it will increase $\rho(d)$ for $2 < d < 3$.

Thus Figure 6 offers a minor insight into spatial network design: that it is city pairs at normalized distance $2-3$ specifically that cause the problems in efficient network design.

The characteristic shape – at least, the flatness over $2 \leq d \leq 5$ – is also visible in the real-world data [10].

For the Hammersley network, the graph of $\rho(d)$ is quite different; $\rho(d)$ increases to a maximum of 0.35 around $d = 0.8$ and then decreases more steeply to a value of 0.21 at $d = 5$. This arises from the particular structure (from each city there is one road in each quadrant) resembling the deterministic "diagonal lattice" of Figure 4, in which the route between some nearby pairs will be via two diagonal roads and a junction.

4 Length-efficiency tradeoff for tractable networks

Recall that our overall theme is the tradeoff between network length and route-length efficiency, and that in this paper we focus on $n \to \infty$ limits in the random model and the particular statistics L and R.

The models described in section 2 are "tractable" in the specific sense that one can find exact analytic formulas for normalized length L. Unfortunately R is not amenable to analytic calculation, and we resort to Monte Carlo simulation to obtain values for R. Table 1 and Figure 7 show the values of (L, R) in the models. We explain below how the values of L are calculated.

Network	L	$\bar{d}$	R
Minimum spanning tree	0.633	2	∞
Relative n'hood	1.02	2.56	0.38
Gabriel	2	4	0.15
Hammersley	3.25	4	0.35
Delaunay	3.40	6	0.07

Table 1. Statistics of tractible networks on random points.

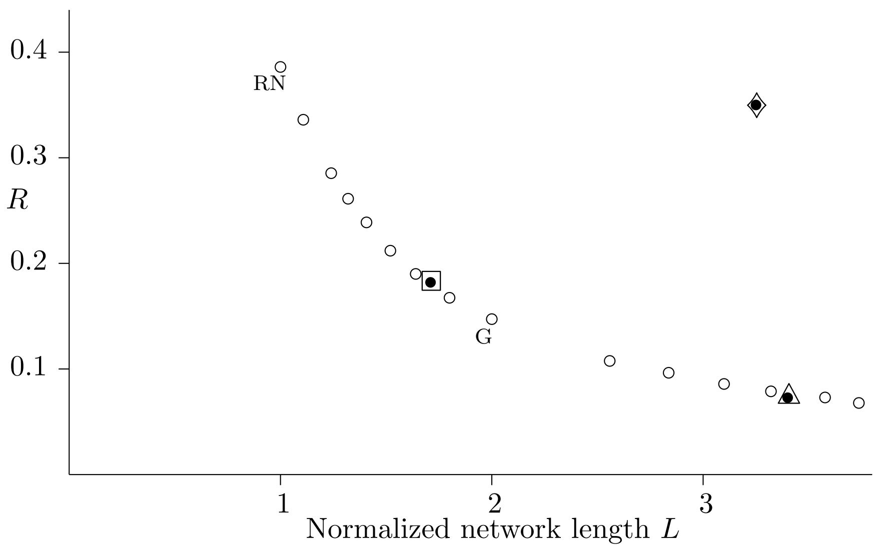

Figure 7. The normalized network length L and the route length efficiency statistic R for certain networks on random points. The $\circ$ show the beta-skeleton family, with RN the relative neighborhood graph and G the Gabriel graph. The $\bullet$ are special models: $\triangle$ shows the Delaunay triangulation, $\square$ shows the network $\mathcal{G}_2$ from section 2.4, and $\diamondsuit$ shows the Hammersley network.

Notes on Table 1. (a) Values of R from our simulations with $n = 2,500$.
(b) Value of L for MST from Monte Carlo [18]. In principle one can calculate arbitrarily close bounds [11] but this has never been carried through. Of course $\bar{d} = 2$ for any tree.
(c) The Gabriel graph and the relative neighborhood graph fit the assumptions of Lemma 1 below with $c = \pi/4$ and $c = \frac{2\pi}{3} - \frac{\sqrt{3}}{4}$ respectively, and their table entries for L and $\bar{d}$ are obtained from Lemma 1, as are the values for β-skeletons in Figure 7.
(d) For the Hammersley network, every degree equals 4, so $L = 2\times$(mean edge-length). It follows from theory [5] that a typical edge, say NE from (x, y), goes to a city at position $(x + \xi_x, y + \xi_y)$, where ξ_x and ξ_y are independent with Exponential(1) distribution. So mean edge-length equals

$$\int_0^\infty \int_0^\infty \sqrt{x^2 + y^2}\ e^{-x-y}\ dxdy \approx 1.62. \tag{7}$$

(e) For any triangulation, $\bar{d} = 6$. For the Delaunay triangulation, $L = ES$ where S is the perimeter length of a typical cell, and it is known ([29] page 113) that $ES = \frac{32}{3\pi}$. Note [28] that the Delaunay triangulation is in general *not* the minimum-length triangulation. Our simulation results in Figure 6 for $\rho(d)$ for the Delaunay triangulation are roughly consistent with a simulation result in [12] saying that $\rho(65) \approx 0.05$.

4.1 A simple calculation for proximity graphs

Lemma 1 *For a proximity graph with template A on the Poisson point process,*

$$L = \frac{\pi^{3/2}}{4c^{3/2}} \quad (8)$$

$$\bar{d} = \frac{\pi}{c} \quad (9)$$

where $c = \text{area}(A)$.

Particular cases can be found elsewhere, (e.g. for $\bar{d}$ in examples in [19]), but our point is to emphasize that Lemma 1 is an elementary and general calculation.

Proof. Take a typical city at position x_0. For a city x at distance s the chance that (x_0, x) is an edge equals $\exp(-cs^2)$ and so

$$\begin{aligned} \text{mean-degree} &= \int_0^\infty \exp(-cs^2)\ 2\pi s\ ds \\ L &= \tfrac{1}{2} \int_0^\infty s \exp(-cs^2)\ 2\pi s\ ds. \end{aligned}$$

Evaluating the integrals gives (9,8).

4.2 Other tractable networks

Curiously, we do not know any other ways of defining networks on random points which are both “natural” and are tractable in the sense that one can find exact analytic formulas for L. In particular we know no tractable way of defining networks with deliberate junctions as in Figure 8 below. Note also that, while it is easy to make ad hoc modifications to the geometric graph to ensure connectivity, these destroy tractability. On the other hand, one can construct “unnatural” networks (see e.g. [8]) designed to permit calculation of L.

5 Optimal networks and $n \to \infty$ limits

5.1 Tractable models

As mentioned earlier, the quantities $L, \bar{d}, R$ we discuss may be interpreted as exact values in the infinite model or as $n \to \infty$ limits in the finite model. To elaborate briefly, in a realization of the finite model (n cities distributed independently and uniformly in a square of area n), a network in Table 1 has a normalized length $L_n = n^{-1} \times$ (network length) and an average degree $\bar{d}_n$ which are random variables, but there is convergence (in probability and in expectation)

$$L_n \to L, \quad \bar{d}_n \to \bar{d} \text{ as } n \to \infty \tag{10}$$

to limit constants definable in terms of the analogous network on the infinite model (rate 1 Poisson point process on the infinite plane). For the proximity graphs or Delaunay triangulation, the network definition applies directly to the infinite model and proof of (10) is straightforward. For the Hammersley network, (10) is implicit in [5], and for the MST a detailed argument can be found in [9]. Convergence of $\rho_n(d)$ to $\rho(d)$ and of R_n to R is similar.

5.2 Optimal networks

We now turn to consideration of *optimal* networks. Given a configuration $\mathbf{x}$ of n cities in the area-n square, and a value of L which is greater than $n^{-1} \times$ (length of Steiner tree), one can define a number

$$R_n(\mathbf{x}, L) = \text{ min. of } \tilde{R} \text{ over all networks on } \mathbf{x} \text{ with normalized length} \leq L \tag{11}$$

where $\tilde{R}$ is the discretized version (6) calculated using intervals of some suitable length δ_n. Applying this to a random configuration $\mathbf{X}$ in the finite model gives, for each L, a random variable

$$\Xi_n(L) := R_n(\mathbf{X}, L).$$

One intuitively expects convergence to some deterministic limit

$$\Xi_n(L) \to R_{\text{opt}}(L), \text{ say, as } n \to \infty. \tag{12}$$

The analogous result for $R_{\max}$ is proved carefully in [8], and the same "super-additivity" argument could be used to prove (12). See [35, 37, 39] for general background to such results. The point is that we don't have any explicit description of the optimal (i.e. attaining the minimum in (11)) networks in the finite or infinite models, so it seems impossible to prove the natural stronger supposition that the finite optimal networks themselves converge (in some appropriate sense) to an infinite optimal network for which the value $R = R_{\mathrm{opt}}(L)$ is attained.

5.3 The curve $R_{\mathrm{opt}}(L)$

Every possible network on the infinite Poisson point process defines a pair (L, R), and the curve $R = R_{\mathrm{opt}}(L)$ can be defied equivalently as the lower boundary of the set of possible values of (L, R). There is no reason to believe that proximity graphs are exactly optimal, and indeed Figure 7 shows that the Delaunay triangulation is slightly more efficient than the corresponding β-skeleton. But our attempts to do better by ad hoc constructions (e.g. by introducing degree-3 junctions – see Figure 8 for an example) have been unsuccessful. We therefore speculate that the function R_{opt} looks something like the curve in Figure 9, which we now discuss.

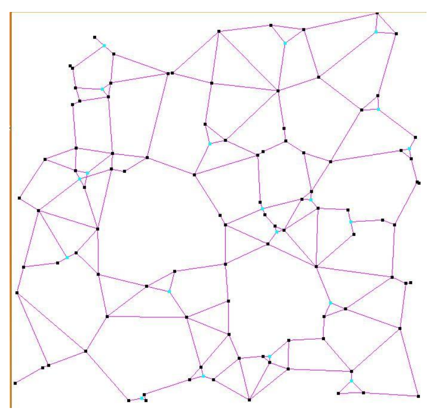

Figure 8. An ad hoc modification of the relative neighborhood graph, introducing junctions.

What can we say about $R_{\mathrm{opt}}(L)$? It is *a priori* non-decreasing. It is known [39] that there exists a *Euclidean Steiner tree constant* L_{ST} representing the limit normalized Steiner tree length in the random model, and clearly $R_{\mathrm{opt}}(L) = \infty$ for $L < L_{\mathrm{ST}}$. The facts

$$R_{\mathrm{opt}}(L) < \infty \text{ for all } L > L_{\mathrm{ST}}; \quad R_{\mathrm{opt}}(L) \to 0 \text{ as } L \to \infty \tag{13}$$

are not trivial to prove rigorously, but follow from the corresponding facts for $R_{\max}$ proved in [8]. But we are unable to prove rigorously that R_{opt} is *strictly* decreasing or that it is continuous.

xxx Quote upper and lower bounds as $L \to \infty$ from final version of [8].

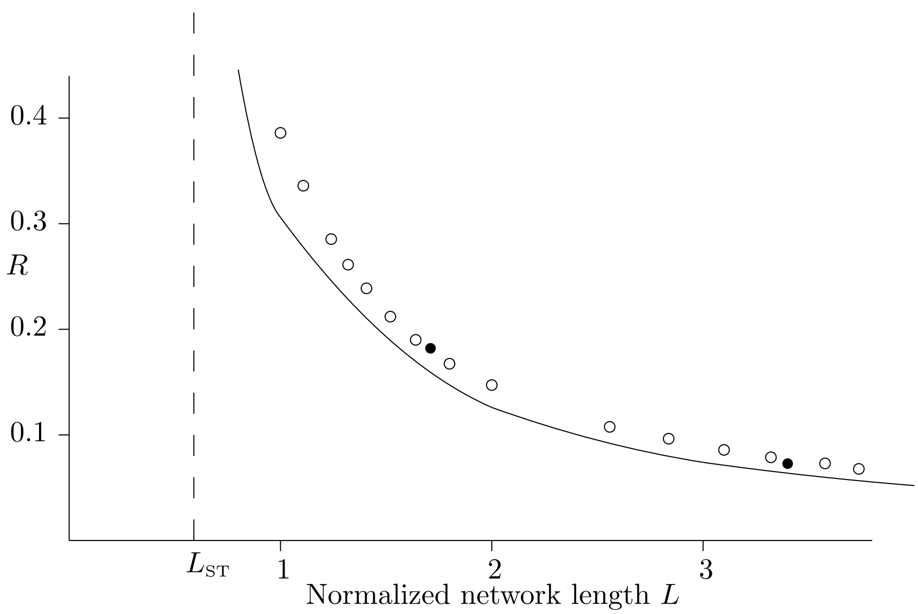

Figure 9. Speculative shape for the curve $R_{\mathrm{opt}}(L)$, with $\circ$ and $\bullet$ values from tractable networks in Figure 7.

6 Final remarks

6.1 Toy models for road networks

The idea of using proximity graphs as toy models for road networks has previously been noted [27] but not investigated very thoroughly. It's intuitively natural to a network designer: whether or not to place a direct road from city i to a nearby city j depends (partly) on whether some other city k is close to the line between them.

For other cost/benefit functionals leading to different optimal networks see [1, 2, 13, 20].

6.2 Rigorous proof of finite R in random proximity graphs?

Table 1 presented the Monte Carlo numerical value ≈ 0.38 of R for the relative neighborhood graph on random points. From a rigorous viewpoint, the assertion that a random network has $R < \infty$ is essentially the assertion that $\rho(d) = O(d)$ as $d \to \infty$. This is often non-trivial to prove. A general sufficient condition for this property, which applies to the relative neighborhood graph (and hence all proximity graphs), is proved in [3]. The related fact that the limit $\lim_{d\to\infty} \rho(d)/d$ exists is proved in [4].

6.3 Optimal trade-off between network length and route-length efficiency

Recall that a central theme is seeking to quantify the trade-off between normalized network length l and route length efficiency R. Figure 9 suggests that for optimal networks the "law of diminishing returns" sets in around $L = 2$ (for comparison, this is the value of L corresponding to the square grid network), in that $R_{\mathrm{opt}}(L)$ decreases rapidly to around 0.13 as L increases to 2 but decreases only slowly as L increases further. We boldly conjecture that for real-world networks in which short route-lengths are a major desideratum, and that are perceived as efficient in actually having short routes, their summary statistics (L, R) will typically be near $(2, 0.2)$. The preliminary data in [10] is roughly consistent with this conjecture.

6.4 Other results for the random network models

There is substantial literature on the networks (MST, proximity graphs, Delaunay triangulation) in the deterministic setting, but the literature for the *random* case is rather diffuse, with different focuses for different networks. For instance, work on MSTs has focused on central limit theorems for network length [26] and connections with critical continuum percolation [16]. For the relative neighborhood graph and the Gabriel graph, [19] calculates $\bar{d}$ and [17] shows that in the finite model, in a certain range the β-skeletons have

$$R_{\mathrm{max}} \text{ grows as order } \sqrt{\log n / \log\log n}. \tag{14}$$

As for the Delaunay triangulation, there has been surprisingly little follow-up to the seminal analysis by Miles [29] (various maximal statistics are studied in [15]), though the closely related Voronoi tessellation has been studied in more detail [30].

6.5 Speculative applications of random proximity graphs

Random proximity graphs seem an interesting object of study from many viewpoints, in particular as an attractive alternative to random geometric graphs for modeling spatial networks that are connected by design. It is remarkable that (14) is the *only* non-elementary result about them that we can find in the literature. As well as being natural models for road networks, they might be useful in modeling communication networks suffering line of sight interference.

At a more mathematical level, for questions such as spread-out percolation [34] or critical value of contact processes [14], random proximity graphs with small A are an interesting alternative to the usual lattice- or random graph-based models. For instance, it is natural to conjecture that the critical value p^*_A for a random proximity graph with template A satisfies

$$p^*_A \sim \pi^{-1}\,\mathrm{area}(A) \text{ as } \mathrm{area}(A) \to 0 \qquad (15)$$

(the right side $= 1/\bar{d}$ from (9)) and that the critical value λ^*_A for the contact process has the same asymptotics.

6.6 Analogies between deterministic and random networks

As mentioned earlier we may make very loose analogies between these networks on random points and particular deterministic networks in Figure 4, based in part on exact equality of $\bar{d}$ in the latter three cases.

Relative n'hood graph	$\leftrightarrow$	punctured lattice
Gabriel graph	$\leftrightarrow$	square lattice
Hammersley network	$\leftrightarrow$	diagonal lattice
Delaunay triangulation	$\leftrightarrow$	triangular lattice

6.7 Scale invariant continuum networks

Introducing the statistic R can be viewed as one approach to resolving the "paradox" from [6], discussed in section 3.2, that the more natural statistic R_{ave} doesn't lead to realistic optimal networks in the $n \to \infty$ limit. This particular approach was prompted by visualizing real-world road networks – cf. discussion in section 3.3. Let us mention a mathematically more sophisticated alternative, under study as work in progress [7]. Instead of a discrete Poisson process of cities we imagine a continuum limit. That is, for each finite set $(z_1, \ldots, z_k)$ of points there is a random network $\mathcal{S}(z_1, \ldots, z_k)$ linking the points, consistent as more points are added. Mathematically natural structural properties for the distribution of such a process are

(i) translation and rotation invariance
(ii) scale invariance
where the latter means that routes, *as point-sets in* $\mathbb{R}^2$, are invariant in distribution under Euclidean scaling. This implies that the quantity $\rho(d)$ analogous to (5), assumed finite, is a constant, which we can call R'. The analog L' of L is defined by

> the expected length of the network on n uniform random points in the area-n square grows $\sim L'n$ as $n \to \infty$.

In this setting we can study the optimal trade-off between L' and R', and the kind of "paradoxical" Figure 5 network cannot arise because it violates scale-invariance.

References

[1] D.J. Aldous. Spatial transportation networks with transfer costs: asymptotic optimality of hub and spoke models *Math. Proc. Cambridge Philos. Soc.*, 145:471–487, 2008.

[2] D.J. Aldous. Optimal spatial transportation networks where link-costs are sublinear in link-capacity. *J. Stat. Mech.*, 2008:P03006, 2008. 12pp.

[3] D.J. Aldous. Which connected spatial networks on random points have linear route-lengths? In preparation, 2009.

[4] D.J. Aldous. The shape theorem for route-lengths in connected spatial networks on random points. In preparation, 2009.

[5] D.J. Aldous and P. Diaconis. Hammersley's interacting particle process and longest increasing subsequences. *Probab. Th. Rel. Fields*, 103:199–213, 1995.

[6] D.J. Aldous and W.S. Kendall. Short-length routes in low-cost networks *via* Poisson line patterns. *Adv. in Appl. Probab.*, 40:1–21, 2008.

[7] D.J. Aldous and W.S. Kendall. Scale-invariant random spatial networks. In preparation, 2009.

[8] D.J. Aldous and T. Lando. The stretch-length tradeoff in geometric networks: Worst-case and average-case study. In preparation, 2009.

[9] D.J. Aldous and J.M. Steele. Asymptotics for Euclidean minimal spanning trees on random points. *Probab. Th. Rel. Fields*, 92:247–258, 1992.

[10] D.J. Aldous and A. Choi. A Route-Length Efficiency Statistic for Road Networks. Unpublished, 2009.

[11] F. Avram and D. Bertsimas. The minimum spanning tree constant in geometric probability and under the independent model: A unified approach. *Ann. Appl. Probab.*, 2:113–130, 1992.

[12] F. Baccelli, K. Tchoumatchenko, and S. Zuyev. Markov paths on the Poisson-Delaunay graph with applications to routeing in mobile networks. *Adv. in Appl. Probab.*, 32:1–18, 2000.

[13] M. Barthélemy and A. Flammini. Optimal traffic networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2006:L07002, 2006.

[14] N. Berger, C. Borgs, J.T. Chayes, and A. Saberi. On the spread of viruses on the internet. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 301–310 (electronic), New York, 2005. ACM.

[15] M. Bern, D. Eppstein, and F. Yao. The expected extremes in a Delaunay triangulation. *Internat. J. Comput. Geom. Appl.*, 1:79–91, 1991.

[16] C. Bezuidenhout, G. Grimmett, and A. Löffler. Percolation and minimal spanning trees. *J. Statist. Phys.*, 92:1–34, 1998.

[17] P. Bose, L. Devroye, W. Evans, and D. Kirkpatrick. On the spanning ratio of Gabriel graphs and β-skeletons. *SIAM J. Discrete Math.*, 20:412–427 (electronic), 2006.

[18] M. Cortina-Borja and T. Robinson. Estimating the asymptotic constants of the total length of Euclidean minimal spanning trees with power-weighted edges. *Statist. Probab. Letters*, 47:125–128, 2000.

[19] L. Devroye. The expected size of some graphs in computational geometry. *Comput. Math. Appl.*, 15(1):53–64, 1988.

[20] M.T. Gastner and M.E.J. Newman. Shape and efficiency in spatial distribution networks. *J. Stat. Mech. Theory Exp.*, pages P01015, 9 pp., (electronic), 2006.

[21] P.-L. George and H. Borouchaki. *Delaunay Triangulation and Meshing.* Editions Hermès, Paris, 1998.

[22] R.L. Graham and P. Hell. On the history of the minimum spanning tree problem. *IEEE Annals of the History of Computing*, 07:43–57, 1985.

[23] A. E. Holroyd, R. Pemantle, Y. Peres, and O. Schramm. Poisson matching. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45:266–287, 2009.

[24] A. E. Holroyd and Y. Peres. Trees and matchings from point processes. *Electron. Comm. Probab.*, 8:17–27 (electronic), 2003.

[25] J.W. Jaromczyk and G.T. Toussaint. Relative neighborhood graphs and their relatives. *Proceedings of the IEEE*, 80(9):1502–1517, 1992.

[26] H. Kesten and S. Lee. The central limit theorem for weighted minimal spanning trees on random points. *Ann. Appl. Probab.*, 6:495–527, 1996.

[27] D.G. Kirkpatrick and J.D. Radke. A framework for computational morphology. In G.T. Toussaint, editor, *Computational Geometry*, pages 217–248. Elsevier, 1985.

[28] E. L. Lloyd. On triangulations of a set of points in the plane. In *18th Annual Symposium on Foundations of Computer Science (Providence, R.I., 1977)*, pages 228–240. IEEE Comput. Soc., Long Beach, Calif., 1977.

[29] R. E. Miles. On the homogeneous planar Poisson point process. *Math. Biosci.*, 6:85–127, 1970.

[30] J. Møller. *Lectures on Random Voronoĭ Tessellations*, volume 87 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 1994.

[31] G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, Cambridge, 2007.

[32] M.E.J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

[33] M. Penrose. *Random Geometric Graphs*. Oxford Univ. Press, 2003.

[34] M. D. Penrose. On the spread-out limit for bond and continuum percolation. *Ann. Appl. Probab.*, 3(1):253–276, 1993.

[35] M. D. Penrose and J. E. Yukich. Weak laws of large numbers in geometric probability. *Ann. Appl. Probab.*, 13:277–303, 2003.

[36] S.Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar. Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol. In *Proc. European Wireless Conference*, 2002.

[37] J.M. Steele. *Probability Theory and Combinatorial Optimization*. Number 69 in CBMS-NSF Regional Conference Series in Applied Math. SIAM, 1997.

[38] D. Stoyan, W. S. Kendall, and J. Mecke. *Stochastic Geometry and its Applications*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons Ltd., Chichester, 2nd edition, 1995. With a foreword by D. G. Kendall.

[39] J.E. Yukich. *Probability Theory of Classical Euclidean Optimization Problems*. Number 1675 in Lecture Notes in Math. Springer, 1998.