Internet Mathematics Vol. 4, No. 2–3: 149–174

On Locality in a Geometric Random Tree Model

Ross M. Richardson

Abstract. We address the question of locality in random graphs. In particular, we study a geometric random tree model $\mathcal{T}_{\alpha,n}$ which is a variant of the FKP model proposed in [Fabrikant et al. 02]. We choose vertices $v_1, \ldots, v_n$ in some convex body uniformly and fix a point $\mathfrak{o}$. We then build our tree inductively, where at time t we add an edge from v_t to the vertex in $v_1, \ldots, v_{t-1}$ that minimizes $\alpha\|v_t - v_i\| + \|v_i - \mathfrak{o}\|$ for $i < t$, where $\alpha > 0$. We categorize an edge $v_i \to v_j$ in this graph as local or global depending on the edge length relative to the distance from v_i to $\mathfrak{o}$. We study the extent to which the tree is composed of either global or local edges and, in particular, show that it undergoes a transition at $\alpha = 1$.

1. Introduction

Consider the problem of providing telephone service to some central hub. Each customer has a given position in the plane, and thus a given distance from the hub. If we are allowed to extend a single connection from each new customer to an existing customer, then choosing the nearest neighbor clearly optimizes (minimizes) the amount of new wire we have to string. On the other hand, connecting to an existing customer who is farther away from the hub than our new customer may lead to attenuated service, and hence we may wish to choose a customer located closer to the hub. If we weigh these two costs and choose a customer who optimizes (minimizes) our total cost function, the behavior will clearly depend on the relative weighting given to each cost.

1542-7951/07 $0.50 per page

We shall construct a simple geometric tree model motivated by the above example. We shall say that a customer is linked by a *local edge* (respectively *global edge*) if the edge length is short (respectively long) *relative to the distance between the customer and the hub.* In this way we obtain a meaningful description of the local behavior that respects the length scale of each vertex. The above example shows that, depending on how we choose to weight edge costs, both completely local and completely global behavior are possible. In this paper, we shall quantify these ideas, and, in particular, we shall look at the transition from global to local behavior based on relative costs.

2. Related Work

Network models that encapsulate both local and global structure have been investigated for some time. One of the earliest such papers is [Bollobás and Chung 88], in which the authors analyze the union of an n-cycle and a random matching. They demonstrate that such graphs have near-minimal diameter (for all graphs of maximum degree at most 3) while at the same time requiring only a linear number of edges. More generally, they also show that a graph with bounded degree k that satisfies an expansion condition (namely that it expands roughly as a tree with constant degree k) can be made to have near-minimal diameter with the addition of a random matching.

From an algorithmic perspective, Kleinberg [Kleinberg 00] proposed a simple local/global network model consisting of an $n \times n$ planar grid, to which one adds a single random edge at each vertex. The edge is chosen with probability proportional to some (fixed) power of the inverse distance. He demonstrates that only for power 2 can an algorithm find short paths given only local information; for other exponents the random component is either too local or too random.

In a similar vein, the authors of [Liben-Nowell et al. 05] show that for an arbitrarily populated grid model in which link probabilities are determined via local density (sparse regions have higher link probability than dense regions), computable short paths exist. The authors of [Andersen et al. 05] replace the local grid with a graph that satisfies prescribed local flow constraints, and the global graph with a power-law $G(\mathbf{w})$ random graph (see [Chung and Lu 06a]). In this way, they obtain both local clustering and small diameter, and they further provide an algorithm for separating the local and global components.

Note that all of these models are constructed as the union of a sparse local component, given by some condition on neighborhood growth (local flow, regular geometric embedding, etc.), combined with a small (linear) number of random edges. The resulting graphs all share the features of being sparse, connected, and having minimal (near-optimal) diameter. As these traits are the hallmark of so-called "real" complex networks, graphs as above with identifiable local and global structure are thus compelling objects of study.

Our main contribution in this work is the analysis of the relative composition of local and global structure in the graph model motivated by the introduction (which is generated similarly to that of [Fabrikant et al. 02]). In particular, we examine the relationship between generative rules (i.e., how should the relative cost between local and global edges be weighted) and the resulting change in overall composition. To our knowledge, the only other work to address the question of detecting local and global structure is [Andersen et al. 05], and they do so algorithmically.

3. Model and Definitions

3.1. Definition

The model we study here is a technical modification of the one originally proposed in [Fabrikant et al. 02] (referred to as the *FKP* model).

We define a random-graph model $\mathcal{T}_\alpha$ ($\mathcal{T}_{\alpha,n}$ when we wish to stress the dependence on n) with positive parameter α. We denote by K some compact, convex set in $\mathbb{R}^d$ and $\mathfrak{o}$ a fixed point interior to K. Without loss of generality, we assume the volume of K to be 1.

For a given natural number n, we construct $\mathcal{T}_{\alpha,n}$ as follows: The vertices of $\mathcal{T}_{\alpha,n}$, denoted by $V_n = \{v_1, \ldots, v_n\}$, are chosen independently and uniformly in K.[1] For each vertex $v_i, i = 1, \ldots, n$, associate to it a function

$$\phi_i(x) = \alpha\|v_i - x\| + \|x - \mathfrak{o}\|,$$

where $\|\cdot\|$ is the Euclidean length. For the model $\mathcal{T}_{\alpha,n}$, we associate to each vertex $v_i, i = 1, \ldots, n$, a unique edge e_i with source v_i and target v_{t_i} given by

$$v_{t_i} := \operatorname*{argmin}_{j<i} \phi_i(v_j).$$

We shall assume that such a minimizer is unique, as this happens with probability 1.

[1] As is usual in the theory of random graphs, we shall adopt the point of view that $\mathcal{T}_{\alpha,n}$ and $\mathcal{T}_{\alpha,m}$ are constructed on the same probability space such that $V_n \cap V_m = V_n$ if $n \leq m$. As the vertices completely determine the graph, the subgraph of $\mathcal{T}_{\alpha,m}$ induced by V_n is thus $\mathcal{T}_{\alpha,n}$, and hence we view $\mathcal{T}$ as being built one vertex and edge at a time.

One obtains the FKP model by choosing K to be the unit square, letting $\mathfrak{o}$ be a vertex of the tree prior to v_1, and using instead the constraints

$$\phi_i'(x) = \alpha\|v_i - x\| + d(x, \mathfrak{o}),$$

where $d(x, \mathfrak{o})$ is the *graph-theoretic* distance from a vertex x to the root $\mathfrak{o}$. Thus the FKP model balances local costs geometrically while balancing global costs graph-theoretically. This definition proves enough of a technical convienience to allow the authors of [Fabrikant et al. 02] to explore the degree distribution of the resulting tree. We summarize these results in the next section.

Let us remark that both our model and the FKP model fall under the paradigm of *highly optimized tolerance* (HOT), due originally to Carlson and Doyle [Carlson and Doyle 00, Carlson and Doyle 99]. So-called HOT models are characterized by "optimal yet reliable design in the presence of a certain hazard." In the present case, our hazard is the randomness introduced by the points in V_n. Through both simulation and comparison with real designed networks, the authors of [Doyle et al. 05] argue that the HOT paradigm more closely reflects the mechanism that drives the formation of such networks, due largely to the fact that real networks evolve (or are designed) to optimize network costs.

3.2. Locality

For an edge e_i with target v_j, we call e_i β-local (or $(1-\beta)$-global) for $0 \leq \beta \leq 1$ if the Euclidean edge length of the orthogonal projection of $\overline{v_i v_j}$ onto the segment $\overline{\mathfrak{o} v_i}$ is at most $\beta\|v_i - \mathfrak{o}\|$. Equivalently, the target of the edge is contained in the half-space

$$H(v_i, \beta) := \left\{x \mid (x - \mathfrak{o}) \cdot (v_i - \mathfrak{o}) \geq (1 - \beta)\|v_i - \mathfrak{o}\|^2\right\}.$$

If not otherwise stated, we take $\beta = \frac{1}{2}$. See Figure 1. We shall concern ourselves with the fraction $\rho(\beta)$ of β-local edges in $\mathcal{T}_\alpha$.

Our definition of a local edge is a straightforward—at least we hope so, given Figure 1—and natural geometric notion. The definition could classify as local an edge emanating from p with length much greater than $\|\mathfrak{o} - p\|$ (which might occur if $\|\mathfrak{o} - p\|$ is small and the edge is perpendicular to $\overline{\mathfrak{o}p}$). In Section 6.1, we introduce the notion of *regions of influence*, which will allow us to quantify the heuristic that the target of p lies "close" to the chord $\overline{\mathfrak{o}p}$ with high probability, and thus abhorrent "local" edges are not a cause for concern. For the present, however, it shall suffice that our definition is a technical convenience.

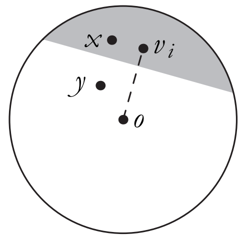

Figure 1. The shaded region represents $H(v_i, 1/4)$, where K is a sphere. Note that x and y (as possible targets of v_i) give 1/4 local and global edges, respectively.

3.3. Conventions

All results hold under the assumption that $n \to \infty$, and we use the Landau notation $O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, etc., with respect to this assumption. Further, we let $\tilde{O}$ suppress logarithmic factors, hence $\tilde{O}(\cdot)$ implies $O\left(\cdot(\log n)^C\right)$ for some constant C. We denote by **P**, **E**, **Var**, **Cov**—the probability, expectation, variance, and covariance, respectively. We say a property holds *asymptotically almost surely* (a.a.s.) if the probability of non-occurrence tends to 0 as $n \to \infty$. The ϵ-ball about a point $p \in \mathbb{R}^d$ is denoted by $B(p, \epsilon)$. We shall use primarily the base-2 logarithm, denoted lg, as well as the natural logarithm, ln.

4. Results

4.1. Prior Results

As $\mathcal{T}_\alpha$ is structurally very similar to the FKP model, it will be useful to understand all that is known in this latter case.

The fundamental result of the authors in [Fabrikant et al. 02] is that the FKP model demonstrates three distinct behaviors based on the value of α.

Theorem 4.1. [Fabrikant et al. 02] *In the FKP model, we have the following:*

- *If $\alpha < 1/\sqrt{2}$, the tree is a star.*
- *If $\alpha = \Omega(\sqrt{n})$, the expected number of nodes that have degree D is at most $n^2 \exp(-cD)$.*
- *If $\alpha \geq 4$ and $\alpha = o(\sqrt{n})$, the expected number of nodes with degree at least D is greater than $c \cdot (D/n)^{-\beta}$ for positive constants c and β (possibly depending on α). In particular, if $\alpha = o\left(\sqrt[3]{n^{1-\varepsilon}}\right)$ then we have $\beta \geq 1/6$ and $c = O(\alpha^{-1/2})$.*

The main achievement of this theorem, namely the appearance of a power law in the model, is disputed to some extent in [Berger et al. 03]. In this later work, the authors take issue with the fact that c in the above theorem may vary with n. While consistent with the work of Fabrikant et al., they show that

- if $\alpha = o(\sqrt{n}/(\log n)^2)$, the FKP model has $n - o(n)$ leaves with high probability, and
- if $\alpha/(\sqrt{n}\log n) \to \infty$, the function ρ_k is bounded above and below by exponential-tailed functions of k,

where $\rho_k = \lim_{n\to\infty} n^{-1}\,\mathbf{E}[|\,\{x \mid x \in V_n, \deg(x) \geq k\}\,|]$.

They further demonstrate that, in the range $(\log n)^4 \leq \alpha \leq n^{1/2}/(\log n)^4$, the maximum degree is $O(n\alpha^{-2})$, yet there are $\tilde{\Theta}\left(\alpha^2\right)$ vertices of degree $\Omega\left(n\alpha^{-2}\right)$. Hence, there are many vertices of degree differing only by a constant from the maximum. This contrasts sharply with the traditional "Zipf-like" notion of a power law, which predicts a constant number of such vertices. They leave unexamined the range for which $\alpha = \Omega\left(\sqrt{n}\,(\log n)^2\right)$ and $\alpha = O\left(\sqrt{n}\log n\right)$. This *logarithmic gap* is "presumably unnecessary" according to the authors, though no more rigorous speculation is given.

As we shall observe in Theorem 4.2, the same dichotomy between "star-like" and exponential-tailed behavior occurs for the model $\mathcal{T}_\alpha$. In this model, however, the behaviors change sharply at $\alpha = 1$. What about the case $\alpha = 1$? Numerical experiments for trees scaling up to 10,000 vertices suggest that in this case the full degree distribution does obey a power law (see Figure 2 for instances of all three regimes). We note that for $\alpha = 1$ the summands in ϕ_i are of the same order, whereas for the FKP model the summands in ϕ_i' are of the same order when $\alpha = \tilde{\Theta}\left(\sqrt{n}\right)$. Thus, for this latter value of $\alpha = \alpha(n)$ we speculate that the FKP model also demonstrates a full power law. Determining the correct value for

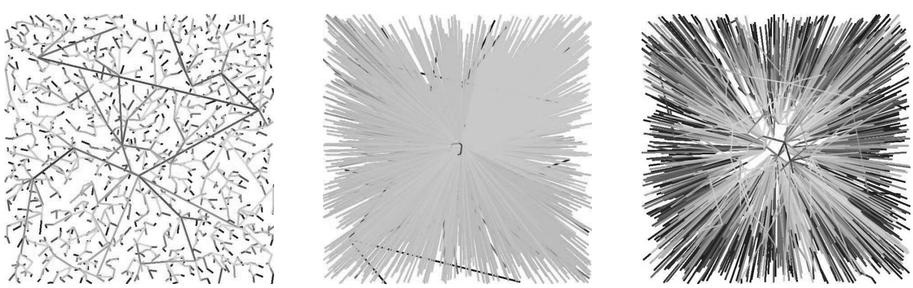

(a) $\alpha = 3$. (b) $\alpha = 0.5$. (c) $\alpha = 1$.

Figure 2. $\mathcal{T}_{\alpha,10{,}000}$ instances.

$\sqrt{n}/(\log n)^2 \leq \alpha \leq \sqrt{n}\log n$ at which this behavior emerges is beyond our reach, however (and hence even convincing numerical computation becomes difficult). Thus, the $\mathcal{T}_\alpha$ model at $\alpha = 1$ seems the correct setting for numerically—if not theoretically—recovering the power law originally sought in [Fabrikant et al. 02].

Finally, in the course of their investigations, the authors also show that the distance from any node to the root $\mathfrak{o}$ is at most $3\log n$ with high probability. They remark that in the case $\alpha = \infty$ the graph-theoretic distance to $\mathfrak{o}$ from a point v_i is distributed as the distance to root in a uniform random-recursive tree (see [Smythe and Mahmoud 95]). Given that the FKP model can be a star with high probability for some α, no logarithmic lower bound on the diameter is possible.

4.2. The Model $\mathcal{T}_\alpha$

Our first two results concern the degree distribution and diameter of $\mathcal{T}_\alpha$ for α bounded away from 1. The model $\mathcal{T}_\alpha$ shares the regimes of "star-like" and exponential-tailed behavior. The transition between regimes, however, occurs here at $\alpha = 1$ instead of $\tilde{\Theta}(\sqrt{n})$ (in the case $d = 2$).

Theorem 4.2. (Degree Distribution.) *We have the following for $\mathcal{T}_{\alpha,n}$:*

1. *If $\alpha < \delta < 1$ for δ fixed, then a.a.s. the number of vertices of degree greater than 1 is $O(\lg n)$.*
2. *If $\alpha > \delta > 1$ for δ fixed, then there exists a constant $c > 0$ such that for any vertex $v_i, i = 1, \ldots, n$, we have*

$$\mathbf{P}[\deg(v_i) \geq D] = O(n\exp(-cD)). \tag{4.1}$$

In particular, the maximum degree is $O(\ln n)$ a.a.s.

Under the assumption that $\alpha \to \infty$ sufficiently fast ($\omega(\ln n)$ say), one can show a matching exponential lower bound in (4.1), which follows by a modification of an argument found in [Berger et al. 03].

Unlike the FKP model, the diameter of our graph is logarithmic for all values of α.

Theorem 4.3. (Diameter.) *For any $\alpha > 0$, the diameter of $\mathcal{T}_\alpha$ is $\Theta(\ln n)$ a.a.s.*

It is worth considering how at the extremes, $\alpha \to 0$ and $\alpha \to \infty$, the diameter is logarithmic. The latter case follows from the prior comparison to a uniform random-recursive tree. The former case, in which the graph appears "star-like", follows from the heuristic that most vertices will choose a target closest to $\mathfrak{o}$ in Euclidean norm. Over the course of n random points, however, the closest point changes a logarithmic number of times. Hence, the diameter is determined by the diameter of the subgraph induced on this sequence of closest points, which will be logarithmic. (This subgraph is expected to be a path since a new "closest" point will target the prior such point.)

Having quantified the combinatorial structure $\mathcal{T}_\alpha$, we now move to the main contribution of this paper: understanding the composition of $\mathcal{T}_\alpha$ in terms of the relative number of local and global edges, as a function of α. Our first result shows that in the case of α bounded away from 1, $\mathcal{T}_\alpha$ consists almost entirely of a single type of edge (local or global).

Theorem 4.4. (Edge Length.) *Fix $0 < \beta < 1$.*

1. *If $\alpha > \delta > 1$, then $\rho(\beta) \to 1$.*
2. *If $\alpha < \delta < 1$, then $\rho(\beta) \to 0$.*

In particular, we find that, when α is bounded away from 1, the distribution of edge lengths is governed primarily by the geometry of the extreme cases, $\alpha \to 0$ and $\alpha \to \infty$. Hence, as $\rho(\beta)$ is a rough measure of the local tendency of the graph, we see that the graph is entirely local or global in this case.

When α is not bounded away from 1, what can be said about $\rho(\beta)$? If β is a function of α, we can have both $\rho(\beta) \to 0$ and $\rho(\beta) \to 1$, and indeed $\rho(\beta)$ can be bounded away from both 0 and 1. While the relationship between α and β that forces these behaviors can be determined to some extent, the attendant details outweigh the utility of such a result. Thus, we focus on our main theorem, which already suggests some of this behavior.

Our main theorem asserts that around $\alpha = 1$, our model $\mathcal{T}_\alpha$ consists of both local and global edges of roughly the same number, as measured by our parameter $\rho(\beta)$. We work in the unit volume ball in $\mathbb{R}^2$, for simplicity, though the same result should hold for general K.

Theorem 4.5. *Set $K = B(0, \pi^{-1/2}) \subseteq \mathbb{R}^2$ with $\mathfrak{o} = 0$. If $\alpha = 1$, we have*

$$\rho\left(\frac{1}{2}\right) = \frac{1}{2}(1 + o(1)), \quad a.a.s.$$

Additionally, if β is fixed and $|\alpha - 1| = o(n^{-2})$, there exist constants $0 < c < c' < 1$ depending on β such that

$$c < \mathbf{E}[\rho(\beta)] < c'.$$

The proof of Theorem 4.5 gives the following heuristic explanation for this behavior: for a point v_i in K with $\|v_i - \mathfrak{o}\| = l$, the edge e_i has length uniformly chosen in $[0, l]$ and is contained in the thin tube about $\overline{\mathfrak{o}v_i}$. It is thus tempting to conjecture that $\rho(x) \approx x$, but the proof of Theorem 4.5 does not seem to generalize to this case.

5. Further Directions

As mentioned in Section 4.1, the model $\mathcal{T}_\alpha$ offers promise as a model on which to study the power-law degree distribution originally sought by [Fabrikant et al. 02]. Of particular interest is the relation between the dimension d and the power-law exponent, which appears in our numerical experiments to be a decreasing function of d.

In both $\mathcal{T}_\alpha$ and the FKP model, the weight α is a fixed function of n. A. Flaxman (personal communication) suggests looking at a non-homogeneous variant of these models. To wit, for the $\mathcal{T}_\alpha$ model, he suggests setting

$$\phi_i(x) = \alpha(i)\|v_i - x\| + \|x - \mathfrak{o}\|,$$

where $\alpha(i)$ varies with i.

The results of Theorem 4.2 and Theorem 4.5 are similar to an infinite Pólya urn model studied in [Chung et al. 03]. In this model, at each time-step a ball is added to an existing urn with probability $1 - p$, else a new urn is created. If a ball is to be placed into an existing urn, then each urn is chosen with probability proportional to m^γ, where m is the number of balls in the urn. Under the regimes $\gamma > 1$, $\gamma < 1$, and $\gamma = 1$, the bin distributions are exponential, dominated by a single bin, and power-law a.a.s. The question of why both combinatorial and geometric selection rules produce similar phenomena remains open.

While random graphs in metric spaces already provide the correct setting for notions of locality, the model $\mathcal{T}_\alpha$ shows that random-graph models with edges chosen in a biased, or spatially inhomogeneous manner, allow for more complex behavior (compare $\mathcal{T}_\alpha$ to the homogeneous geometric random-graph model of Penrose). One outstanding question available in this setting is understanding the extent to which dimensionality can be defined intrinsically for complex network models, and how dimensionality influences the degree distribution, edge locality, etc. Such a program would be an important first step in the verification of geometric random-graph models.

6. Tools

6.1. Regions of Influence

All of the subsequent analysis in this paper relies on the notion of *influence regions*. Set $\gamma > 0$. Then the set of points

$$U(p,\gamma) = \{q \mid \|\mathfrak{o} - q\| + \alpha\|q - p\| \le (\min(1,\alpha)+\gamma)\|p - \mathfrak{o}\|\}$$

forms the γ-*influence region* (or $\gamma-$region) about p. We then have the following:

Lemma 6.1. (Convexity.) *The region $U(p,\gamma)$ is convex for any choice of p and $\gamma > 0$.*

Proof. Let $\mathbf{x}_1$, $\mathbf{x}_2 \in K$ be such that

$$\|\mathbf{x}_i - \mathfrak{o}\| + \alpha\|\mathbf{x}_i - p\| = (\min(1,\alpha)+\gamma)\|p - \mathfrak{o}\|, \quad i = 1,2, \tag{6.1}$$

which is to say that they lie on the boundary of $U(p,\gamma)$. Let $\mathbf{z} = \lambda\mathbf{x}_1 + (1-\lambda)\mathbf{x}_2$. Then

$$\|\mathbf{z} - p\| = \|\lambda\mathbf{x}_1 + (1-\lambda)\mathbf{x}_2 - p\| \le \lambda\|\mathbf{x}_1 - p\| + (1-\lambda)\|\mathbf{x}_2 - p\|,$$

and similarly for $\|\mathbf{z} - \mathfrak{o}\|$. Thus,

$$\begin{aligned}\|\mathbf{z} - \mathfrak{o}\| + \alpha\|\mathbf{z} - p\| &\le (\lambda\|\mathbf{x}_1 - \mathfrak{o}\| + (1-\lambda)\|\mathbf{x}_2 - \mathfrak{o}\|)\\ &\quad + \alpha(\lambda\|\mathbf{x}_1 - p\| + (1-\lambda)\|\mathbf{x}_2 - p\|)\\ \text{(by (6.1))} &= (\min(1,\alpha)+\gamma)\|p - \mathfrak{o}\|.\end{aligned}$$

□

For the special case $\alpha = 1$, the γ-region is simply an ellipse with foci $\mathfrak{o}$ and p and with major axis length $\|\mathfrak{o} - p\|(1+\gamma)/2$. For $\alpha < 1$, as $\gamma \to \infty$ the γ region approaches that of an ellipse with foci $\mathfrak{o}$ and p. For γ sufficiently small, however, the γ-region localizes about the point $\mathfrak{o}$. Specifically, the γ-region forms a convex region about $\mathfrak{o}$, the boundary of which is at maximum and minimum distance from $\mathfrak{o}$ along the line through $\mathfrak{o}$ and p. The case $\alpha > 1$ is similar, but in this case the γ-region concentrates about p. See Figure 3.

We can further elucidate the structure of this region by computing the radii of the smallest enclosing circle and the largest inscribing circle of $U(p,\gamma)$. We summarize this as follows:

Lemma 6.2. *Let p be a point of distance r to $\mathfrak{o}$.*

Assume $\alpha < 1$. Then we have

$$B\left(\mathfrak{o}, \frac{r(1+\gamma-\alpha)}{1+\alpha}\right) \subset U(p,\gamma) \subset B\left(\mathfrak{o}, \frac{r(1+\gamma-\alpha)}{1-\alpha}\right),$$

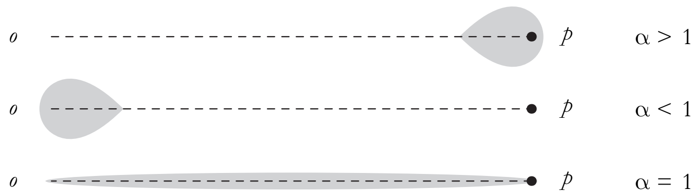

Figure 3. $\gamma-$regions for γ small.

and, in particular,

$$\frac{\pi_d}{(1+\alpha)^d} \le \frac{\text{Area}(U(p,\gamma))}{(r(1+\gamma-\alpha))^d} \le \frac{\pi_d}{(1-\alpha)^d}.$$

Let $\alpha > 1$. *Then we have*

$$B\left(\mathfrak{o}, \frac{r\gamma}{1+\alpha}\right) \subset U(p,\gamma) \subset B\left(\mathfrak{o}, \frac{r\gamma}{1-\alpha}\right),$$

and

$$\frac{\pi_d}{(1+\alpha)^d} \le \frac{\text{Area}(U(p,\gamma))}{(r\gamma)^d} \le \frac{\pi_d}{(1-\alpha)^d}.$$

Here, $\pi_d = \frac{\pi^{d/2}}{\Gamma(d/2+1)}$ denotes the volume of a unit ball in $\mathbb{R}^d$.

Proof. For $\alpha > 1$, consider the ball centered at p of minimum radius xr, $x \le 1$, that includes $U(p,\gamma)$. As $U(p,\gamma)$ and this ball intersect along the line $p\mathfrak{o}$, we obtain the equation $(1-x)r + \alpha xr = (1+\gamma)r$, hence $xr = \gamma r/(1-\alpha)$. The other cases are similar. □

It is worth observing that, whenever $\alpha \neq 1$, the ratio of the radius of a ball that inscribes $U(p,\gamma)$ to the radius of a ball that circumscribes $U(p,\gamma)$ can be made at most $\frac{\alpha+1}{|\alpha-1|}$ for any γ. As a consequence, for α bounded away from 1, we can conclude that the length of e_i is of the same order as the nearest neighbor distance of v_i. See Lemmas 7.1 and 7.5.

6.2. Probability

The following elementary inequality will be used without comment in this paper:

$$(1-nx^2)e^{-nx} \le (1-x)^n \le e^{-nx}, \quad 0 \le x \le 1, \quad n \ge 0.$$

See for instance [Bai et al. 01], Lemma 5.

We shall also make use of the following Chernoff-like bound for sums of independent indicators (see [Chung and Lu 06b] for a recent survey).

Lemma 6.3. **(Chernoff.)** *Let $X_1, \ldots, X_n$ be independent random variables such that $\mathbf{P}(X_i = 1) = p_i$ and $\mathbf{P}(X_i = 0) = 1 - p_i$. Then if $X = \sum_{i=1}^n X_i$, we have the following two bounds:*

$$\mathbf{P}(X \leq \mathbf{E}[X] - \lambda) \leq \exp\left(\frac{-\lambda^2}{2\,\mathbf{E}[X]}\right),$$
$$\mathbf{P}(X \geq \mathbf{E}[X] + \lambda) \leq \exp\left(\frac{-\lambda^2}{2(\mathbf{E}[X] + \lambda/3)}\right).$$

The next lemma is a convenient way to invoke a second-moment argument, following [Alon and Spencer 00], and follows from Chebyshev's inequality.

Lemma 6.4. *Let $X = X_1 + \ldots + X_n$ where X_i is an indicator for the event A_i. If $\mathbf{Var}[X] = o(\mathbf{E}[X]^2)$, then we have*

$$X = \mathbf{E}[X](1 + o(1)), \quad a.a.s.$$

7. Proofs

7.1. Proof of Theorem 4.2

7.1.1. Proof of case $\alpha < \delta < 1$. For $i = 1, \ldots, \lfloor \lg n \rfloor$ we let $r_i = (\frac{2i}{\pi_d (\lg e)(2^i - 1)})^{1/d}$. We claim that

$$\mathbf{P}\left(V_{2^i-1} \cap B(\mathfrak{o}, r_i) = \emptyset \text{ for some } i = \lfloor \lg\lg n \rfloor, \ldots, \lfloor \lg n \rfloor\right) = o(1).$$

To see this, we use the union bound to compute

$$\begin{aligned} \mathbf{P}\left(\bigcup_{i=\lfloor \lg\lg n \rfloor}^{\lfloor \lg n \rfloor} \{V_{2^i-1} \cap B(\mathfrak{o}, r_i) = \emptyset\}\right) &\leq \sum_{i=\lfloor \lg\lg n \rfloor}^{\lfloor \lg n \rfloor} (1 - \pi_d r_i^d)^{2^i-1} \\ &\leq \sum_{i=\lfloor \lg\lg n \rfloor}^{\lfloor \lg n \rfloor} \exp(-\pi_d r_i^d (2^i - 1)) \\ &\leq \lg n (\lg n)^{-3/2} = o(1), \end{aligned} \tag{7.1}$$

since $\exp(-2\lfloor \lg\lg n \rfloor / \lg e) \leq (\lg n)^{-3/2}$ for n sufficiently great.

Next, observe that the point closest to $\mathfrak{o}$ gives a bound on the target for any vertex.

Lemma 7.1. *Fix* $2 \leq i \leq n$. *If* $r = \min_{j<i} \|\mathfrak{o} - v_j\|$, *then* $\|\mathfrak{o} - v_{t_i}\| < r\left(\frac{1+\delta}{1-\delta}\right)$.

Proof. Let v_j be the point that obtains the minimum distance r. By Lemma 6.2, there is a γ such that $B(\mathfrak{o}, r) \subset U(v_i, \gamma) \subset B(\mathfrak{o}, r\frac{1+\alpha}{1-\alpha})$. Thus, any point v_k falling outside $B(\mathfrak{o}, r\frac{\alpha+1}{\alpha-1})$ must have $\phi_i(v_k) > \phi_i(v_j)$, and hence the target of v_i must obey $\|\mathfrak{o} - v_{t_i}\| \leq r\frac{1+\alpha}{1-\alpha} < r\frac{1+\delta}{1-\delta}$. □

Now, for each $i = \lfloor \lg\lg n \rfloor, \ldots, \lfloor \lg n \rfloor$, we can compute the number of points in $\{v_{2^i}, \ldots, v_{2^{i+1}-1}\}$ that fall inside $B(\mathfrak{o}, r_i(\frac{1+\delta}{1-\delta}))$, which for n sufficiently great lies entirely inside K. The expected such number is $2^i\pi_d(r_i\frac{1+\delta}{1-\delta})^d = \frac{2i}{\lg e}(\frac{2^i}{2^i-1})(\frac{1+\delta}{1-\delta})^d$. By Lemma 6.3, the probability that the actual number exceeds the expectation by more than ci for some constant $c > 0$ is bounded by

$$\exp\left(\frac{-c^2i^2}{2\left(\frac{2i}{\lg e}\left(\frac{2^i}{2^i-1}\right)\left(\frac{1+\delta}{1-\delta}\right)^d + ci/3\right)}\right),$$

which is $o((\lg n)^{-1})$ for c sufficiently great and $i \geq \lfloor \lg\lg n \rfloor$. Hence, we may assume that for all such i the number of points in $\{v_{2^i}, \ldots, v_{2^{i+1}-1}\}$ falling inside $B(\mathfrak{o}, r_i\frac{1+\delta}{1-\delta})$ is $\frac{2i}{\lg e}(\frac{2^i}{2^i-1})(\frac{1+\delta}{1-\delta})^d + ci$ with probability tending to 1. Summing gives

$$\sum_{i=\lfloor \lg\lg n \rfloor}^{\lfloor \lg n \rfloor} i\left(\left(\frac{2^i}{2^i-1}\right)\left(\frac{1+\delta}{1-\delta}\right)^d + c\right) = O(\lg n)$$

such vertices. From Equation (7.1) and Lemma 7.1 the remaining vertices among $\{v_{2^{\lfloor \lg\lg n \rfloor}}, \ldots, v_n\}$ have degree 1. Hence, the total number of vertices of degree greater than 1 is $O(\lg n)$.

7.1.2. Proof of case $\alpha > \delta > 1$. The method of proof here is an extension of that found in [Fabrikant et al. 02].

Fix the point v_i. We consider each edge in which v_i participates to be either *short* or *long* based on whether or not it is shorter than $r = (\pi_d(n-1))^{-1/d}$. The number of short edges of v_i is bounded by the number of vertices that fall inside the ball of radius r. Lemma 6.3 shows that

$$\mathbf{P}(\#\text{short edges } \geq D/2) \leq \exp\left(\frac{(D/2-1)^2}{4/3 + D/3}\right).$$

Now, pick ϵ and θ sufficiently small such that

$$\frac{(1+\epsilon)}{\delta} + \sqrt{2((1+\epsilon)^2 - \cos\theta)} < 1, \tag{7.2}$$

observing that the left-hand side decreases in both ϵ and θ. We then have the following lemma.

Lemma 7.2. *For $l > 0$, let v_j, v_k lie in the annulus about v_i of radii l and $(1+\epsilon)l$ for any $l > 0$. If $\measuredangle v_j v_i v_k < \theta$, then $\phi_j(v_k) < \phi_j(v_i)$; i.e., v_j prefers v_k to v_i.*

Proof. We abbreviate $d_{xy} := \|x - y\|$. By the law of cosines,

$$d^2_{v_j v_k} \leq d^2_{v_j v_i} + d^2_{v_k v_i} - 2 d_{v_j v_i} d_{v_k v_i} \cos \measuredangle v_j v_i v_k \leq 2(1+\epsilon)^2 l^2 - 2l^2 \cos\theta$$

hence $d_{v_j v_k} \leq l\sqrt{2((1+\epsilon)^2 - \cos\theta)}$. Further, $d_{v_i v_j} \geq l$. Hence, by (7.2)

$$(1+\epsilon)l < \delta(l - l\sqrt{2((1+\epsilon)^2 - \cos\theta)}) \leq \delta(d_{v_i v_j} - d_{v_j v_k}) < \alpha(d_{v_i v_j} - d_{v_j v_k}).$$

On the other hand, $d_{v_i \mathfrak{o}} - d_{v_j \mathfrak{o}} \leq (1+\epsilon)l$, so $d_{v_k \mathfrak{o}} - d_{v_i \mathfrak{o}} < \alpha(d_{v_i v_j} - d_{v_j v_k})$; i.e.,

$$d_{v_k \mathfrak{o}} + \alpha d_{v_j v_k} < d_{v_i \mathfrak{o}} + \alpha d_{v_i v_j},$$

our conclusion. □

In $\mathbb{R}^d$, no more than $N_0 = N_0(d, \theta)$ points can be placed such that the angle $\measuredangle x v_i y > \theta$ for every pair x, y. This follows by a simple packing argument. As a result, in any annulus as above there can be at most N_0 vertices linked to v_i. Now, if L_x denotes the number of vertices linked to v_i in the annulus of radii x and $(1+\epsilon)x$ centered at v_i, we can thus count the total number of long edges as

$$\sum_{i=\lfloor \log_{1+\epsilon} \zeta_i \rfloor}^{\lfloor \log_{1+\epsilon} \mathrm{diam} K \rfloor} L_{(1+\epsilon)^i} \leq N_0(\log_{1+\epsilon} \mathrm{diam} K - \log_{1+\epsilon} \zeta_i + 1),$$

where $\zeta_i = \min_j \{\|v_i - v_j\|, r\}$. Now, $\mathbf{P}(N_0(\log_{1+\epsilon} \mathrm{diam} K - \log_{1+\epsilon} \zeta_i + 1) \geq D/2)$ is just $\mathbf{P}(\zeta_i \leq \mathrm{diam} K (1+\epsilon)^{1-D/2N_0})$. As our points are chosen independently, the union bound gives $\mathbf{P}(\zeta_i \leq \mathrm{diam} K(1+\epsilon)^{1-D/2N_0}) \leq (n-1)\pi_d(\mathrm{diam} K(1+\epsilon)^{1-D/2N_0})^d$. Thus,

$$\mathbf{P}(\#\text{long edges } \geq D/2) = O(n \exp(-cD))$$

for some constant c.

As $\mathbf{P}(d(v_i) \geq D) \leq \mathbf{P}(\#\text{short edges } \geq D/2) + \mathbf{P}(\#\text{long edges } \geq D/2)$, we thus have $\mathbf{P}(d(v_i) \geq D) = O(n\exp(-c'D))$ for some $c' > 0$.

Finally, note that if we set $\Delta_0 = \frac{3\ln n}{c'}$ (here 3 can be replaced with $2+\varepsilon$), we have that $\mathbf{P}(\deg(v_i) \geq \Delta_0) = O(n^{-2})$. Hence, by the union bound, we have

$$\mathbf{P}(\Delta(\mathcal{T}_\alpha) \geq \Delta_0) \leq \sum_{i=1}^{n} \mathbf{P}(\deg(v_i) \geq \Delta_0) = nO(n^{-2}) = O(n^{-1}),$$

so the maximum degree of $\mathcal{T}_\alpha$ is bounded by Δ_0 a.a.s.

7.2. Proof of Theorem 4.3

The upper bound follows from a simple adaptation of 2 in [Berger et al. 03], so we focus on the lower bound. We shall use the following, which is adapted from Lemma 1 in [Berger et al. 03].

Lemma 7.3. *The probability that, for* $1 < i_1 < \ldots < i_k$*, the path* $v_{i_k} \to v_{i_{k-1}} \to \ldots \to v_{i_1} \to v_1$ *exists in* $\mathcal{T}_\alpha$ *is* $1/((i_1-1)(i_2-1)\cdots(i_{k-1}-1))$.

Proof. We shall fix the positions—but not the labels—of vertices in V_{i_k-1}. As edges are determined only from the positions, the target of v_{i_k} is any member of V_{i_k-1}, chosen uniformly with probability $1/(i_k-1)$. If we condition on the target being $v_{i_{k-1}}$, we may view the vertex labels prior to $v_{i_{k-1}}$ as still unchosen, and thus we may repeat our argument to obtain the lemma. □

To prove our lower bound, we bound the length of the $v_n \to v_1$ path in $\mathcal{T}_\alpha$. Note that the probability that the length of this path is at most L is given by

$$\frac{1}{n-1} + \sum_{r=2}^{L} \sum_{1<i_1<\ldots<i_{r-1}<n} \frac{1}{(i_1-1)\cdots(i_{r-1}-1)(n-1)} \leq \frac{1}{n-1}\sum_{r=1}^{L}\frac{(1+\ln n)^{r-1}}{r!}. \quad (7.3)$$

Now, if $a_r = \frac{(1+\ln n)^{r-1}}{r!}$, then $a_{r+1} = a_r\frac{1+\ln n}{r+1}$, and hence for a given n the sequence $\{a_r\}_{r=1}^{\infty}$ is unimodal with maximum $a_{\lceil \ln n\rceil}$. If we let $L = \lfloor \ln n/2 \rfloor$, then Stirling's approximation shows a_L (and hence all lower terms) to be $o(n/\ln n)$, and hence (7.3) tends to zero. Thus, with probability tending to 1 the path from $v_n \to v_1$ is greater than $L = \lfloor \ln n/2 \rfloor$.

7.3. Proof of Theorem 4.4

We shall handle only the case $\alpha > \delta > 1$, as the other case is similar.

We will require a small technical fact, which is a (more or less) direct consequence of the compactness of K.

Lemma 7.4. *Let $x \in K$. Then there exists an $r_0 > 0$ and $\eta > 0$ such that* $\mathrm{Vol}(K \cap B(x,r)) \geq \eta\pi_d r^d$ *for all $r \leq r_0$.*

Now, for each $i = 2, \dots, n$, we shall set $r_i = (\frac{2\ln(i-1)}{\eta\pi_d(i-1)})^{1/d}$. The probability that all vertices in V_{i-1} are a distance at least r_i from v_i is $(1-\mathrm{Vol}(B(v_i,r_i)\cap K))^{i-1} \leq (1-\eta\pi_d r^d)^{i-1} \leq \exp(-2\ln(i-1)) = (i-1)^{-2}$. Hence, we may assume that for each $i \geq i_0$, $V_{i-1} \cap B(v_i, r_i) \neq \emptyset$ with probability $\sum_{i=i_0}^n (i-1)^{-2}$, which tends to zero as $i_0 = i_0(n) \to \infty$.

The following fact is a companion to Lemma 7.1 and is proved in the same way.

Lemma 7.5. *Fix $2 \leq i \leq n$. If $r = \min_{j<i} \|v_i - v_j\|$ then $\|v_i - v_{t_i}\| < r\left(\frac{\delta+1}{\delta-1}\right)$.*

In particular, note that if $B(v_i,r_i) \cap V_{i-1} \neq \emptyset$ then $\|v_i - v_{t_i}\| < r_i(\frac{\delta+1}{\delta-1})$. If, in addition, $\|v_i - \mathfrak{o}\| > \beta^{-1} r_i(\frac{\delta+1}{\delta-1})$ then e_i is β-local. The probability that v_i falls in $B(\mathfrak{o}, \beta^{-1} r_i(\frac{\delta+1}{\delta-1}))$ is bounded by $\frac{\beta^{-d}\ln(i-1)}{\eta(i-1)}$. Thus, the expected number of such "close" points is at most $\frac{\beta^{-d}}{\eta}\sum_{i=1}^n \frac{\ln(i-1)}{i-1} = O((\ln n)^2)$. By Lemma 6.3, the actual number of such points is $O((\ln n)^2)$ with probability tending to 1 (taking $\lambda = c(\ln n)^{2/3}$, say, for c sufficiently great). Thus, the only edges that can fail to be β-local are these "close" points as well as the first i_0. Setting $i_0 = O((\ln n)^2)$ thus allows us to conclude that the number of local edges is $n(1 - o(1))$ with probability tending to 1.

7.4. Proof of Theorem 4.5, $\beta = 1/2$

Our argument will be via the second-moment method, using Lemma 6.4. The key idea is that, for $\alpha = 1$, the γ-region for each point $v_i \in V_n$ is (with high probability) an *ellipse*. In particular, reflecting all points except v_i in such an ellipse about the minor axis does not alter their their values with respect to ϕ_i. Thus, the probability that any given edge is local is about $\frac{1}{2}$. Unfortunately, as two γ-regions always overlap (they always contain the point $\mathfrak{o}$, for example), we must overcome issues of dependency in applying Lemma 6.4.

We shall let X_i be the indicator that vertex v_i is local, hence $X = \sum_{i=1}^n X_i$ counts the number of local vertices. For a given point $x \in K$, we shall let $E(x,t)$ denote the (closed) ellipse with foci x and $\mathfrak{o}$ and area t.

We shall partition the vertices V_n into *epochs*, indexed by $0, 1, \ldots, \lfloor \lg n \rfloor$, where vertex v_i belongs to epoch $l(i) := \lfloor \lg i \rfloor$. To each epoch we associate a set of parameters that are given (with foresight) by the following:

- $t_i = 2^{-i/2}$,
- $r_i = \left(1 - 2^{-1/7}\right)^{i/2}$,
- $\theta_i = \left(1 - 2^{-1/7}\right)^{i}$,
- and $h_i = \left(\frac{2-\varepsilon}{4}\right)^{i/2}$ for some $0 < \varepsilon < 2 - \frac{1}{1-2^{-1/7}}$.

Given these parameters, we shall construct the following events:

- $\mathcal{A}(v_i, l)$: $E(v_i, t_l) \subseteq K$,
- $\mathcal{B}(v_i, l)$: $|(E(v_i, t_l) - B(\mathfrak{o}, h_l) - B(v_i, h_l)) \cap V_i| > 1$, $\|v_i - \mathfrak{o}\| > r_l$, $E(v_i, t_l) \cap B(\mathfrak{o}, h_l) \cap V_{i-1} = \emptyset$, and $E(v_i, t_l) \cap B(v_i, h_l) \cap V_{i-1} = \emptyset$,
- $\mathcal{C}(v_i, v_j, l)$: $\measuredangle v_i \mathfrak{o} v_j > \theta_l$.

The following geometric lemma will be of critical importance.

Lemma 7.6. *If $t = o(r^2)$ and $r \to 0$, then there exists a constant $c > 0$ such that the distance from $\mathfrak{o}$ to any point in $E(v_i, t) \cap E(v_j, t)$ is bounded by $c\frac{t^2}{r^5\left(1-\cos\frac{\theta}{2}\right)}$.*

Finally, when examining the covariance of two indicators X_i and X_j, we will need the (high probability event) that

1. neither v_i or v_j lies too close to the boundary of K nor too close to $\mathfrak{o}$,
2. the ellipses $E(v_i, t_l)$ and $E(v_j, t_l)$ each contain points from $V_n \setminus \{v_i, v_j\}$ which do not lie near the ends of the major axes, and
3. v_i and v_j form a sufficiently great angle with $\mathfrak{o}$.

We define this event $\mathcal{D}$ according to our above notation as

$$\mathcal{D}(v_i, v_j, l) = \mathcal{A}(v_i, l) \cap \mathcal{A}(v_j, l) \cap \mathcal{B}(v_i, l) \cap \mathcal{B}(v_j, l) \cap \mathcal{C}(v_i, v_j, l).$$

We then have the following technical lemma.

Lemma 7.7. *The following estimates hold:*

$$\sum_{i=1}^{n} \mathbf{P}(\overline{\mathcal{A}(v_i, l(i) \cap \mathcal{B}(v_i, l(i)))}) = o(n), \quad (7.4)$$

$$\sum_{i=1}^{n} \sum_{j>i}^{n} \mathbf{P}(\overline{\mathcal{D}(v_i, v_j, l(j))}) = o(n^2), \quad (7.5)$$

and

$$\sum_{i=1}^{n} \sum_{j>i}^{n} \mathbf{P}(\overline{\mathcal{A}(v_i, l(j)) \cap \mathcal{B}(v_i, l(j))}) + \mathbf{P}(\overline{\mathcal{A}(v_j, l(j)) \cap \mathcal{B}(v_j, l(j))}) = o(n^2). \quad (7.6)$$

Here, the $\overline{\mathcal{E}}$ denotes the complement of an event $\mathcal{E}$. We can now prove Theorem 4.5.

Proof of Theorem 4.5 (Part I). *Expectation.* Fix $1 \leq i \leq n$. We shall construct a map T^i of the underlying probability space K^n. On the complement of the set $\mathcal{A}(v_i, l(i)) \cap \mathcal{B}(v_i, l(i))$, T^i acts as the identity. Otherwise, for each vertex $v_j \neq v_i$ the map T^i is the identity if $v_j \notin E(v_i, t_{l(i)})$ and otherwise reflects v_j about the minor axis of $E(v_i, t_{l(i)})$. The assumption $\mathcal{A}(v_i, l(i))$ makes this operation well defined. Further, T^i is probability-preserving as rigid reflection preserves Lebesgue measure.

Under $\mathcal{B}(v_i, l(i))$, the target of v_i lies inside $E(v_i, t_{l(i)}) - B(\mathfrak{o}, h_{l(i)}) - B(v_i, h_{l(i)})$. As $E(v_i, t_{l(i)}) - B(\mathfrak{o}, h_{l(i)}) - B(v_i, h_{l(i)})$ is symmetric about the minor axis of $E(v_i, t_{l(i)})$, T^i exchanges the events $X_i = 1$ and $X_i = 0$ under $\mathcal{A}(v_i, l(i)) \cap \mathcal{B}(v_i, l(i))$. Hence, we have

$$\mathbf{P}(X_i = 1) = 1/2 + O\left(1 - \mathbf{P}\left(\mathcal{A}(v_i, l(i)) \cap \mathcal{B}(v_i, l(i))\right)\right).$$

Thus, we can compute

$$\mathbf{E}[X] = \sum_{i=1}^{n} 1/2 + O\left(1 - \mathbf{P}\left(\mathcal{A}(v_i, l(i)) \cap \mathcal{B}(v_i, l(i))\right)\right) = n/2 + o(n), \quad (7.7)$$

where the last equality follows from (7.4).

Variance. Computation of the variance is similar. We shall fix $1 \leq i < j \leq n$. We now construct two maps S_l^i on K^n for each l. The map S_l^i fixes all those v_j that fall outside $E(v_i, t_l)$ as well as v_i. Those vertices that fall inside $E(v_i, t_l)$ are reflected about the minor axis of this ellipse.

If we set $l_0 = \lg \lg n$, we can see that for $l \geq l_0$ we have $r_l \to 0$ and $t_l = o(r_l^2)$. Hence, by our choice of the parameters t_l, r_l, h_l, and θ_l as well as Lemma 7.6 we see that on the event $\mathcal{D}(v_i, v_j, l(j))$ the regions $E(v_i, t_{l(j)}) - B(v_i, h_{l(j)}) - B(\mathfrak{o}, h_{l(j)})$ and $E(v_j, t_{l(j)}) - B(v_j, h_{l(j)}) - B(\mathfrak{o}, h_{l(j)})$ are disjoint as long as $l(j) \geq l_0$, and hence $S^i_{l(j)}$ and $S^j_{l(j)}$ commute. Further, both these maps are probability-preserving on $\mathcal{D}(v_i, v_j, l(j))$, so we see that the events $\{X_i = \varepsilon_1 \wedge X_j = \varepsilon_2\} \cap \mathcal{D}(v_i, v_j, l(j))$ consist of the orbit of any one of them under the group generated by $S^i_{l(j)}$ and $S^j_{l(j)}$ and hence are equiprobable. Thus, $\mathbf{E}[X_i X_j] \leq 1/4 + \mathbf{P}\left(\overline{\mathcal{D}(v_i, v_j, l(j))}\right)$. As a result, we have that

$$\begin{aligned}
\mathbf{Cov}[X_i X_j] &= \mathbf{E}[X_i X_j] - \mathbf{E}[X_i]\,\mathbf{E}[X_j] \\
&\leq 1/4 + \mathbf{P}\left(\overline{\mathcal{D}(v_i, v_j, l(j))}\right) \\
&\quad - \left(1/2 - 1/2\,\mathbf{P}(\overline{\mathcal{A}(v_i, l(j)) \cap \mathcal{B}(v_i, l(j))})\right) \\
&\quad \times \left(1/2 - 1/2\,\mathbf{P}(\overline{\mathcal{A}(v_j, l(j)) \cap \mathcal{B}(v_j, l(j))})\right) \\
&\leq \mathbf{P}\left(\overline{\mathcal{D}(v_i, v_j, l(j))}\right) \\
&\quad + 1/4\left(\mathbf{P}(\overline{\mathcal{A}(v_i, l(j)) \cap \mathcal{B}(v_i, l(j))}) + \mathbf{P}(\overline{\mathcal{A}(v_j, l(j)) \cap \mathcal{B}(v_j, l(j))})\right).
\end{aligned} \tag{7.8}$$

Hence, we can compute

$$\begin{aligned}
\mathbf{Var}[X] &= \sum_{i=1}^{n} \mathbf{Var}[X_i] + 2\sum_{i=1}^{n}\sum_{j>i}^{n} \mathbf{Cov}[X_i X_j] \\
&\leq n + 2\sum_{i=1}^{n}\sum_{j>i} \mathbf{Cov}[X_i X_j] \\
\text{(by (7.8))} &\leq n + 2\sum_{i=1}^{n}\sum_{j>i}^{n} \Bigg(\mathbf{P}\left(\overline{\mathcal{D}(v_i, v_j, l(j))}\right) \\
&\qquad + 1/4\Big(\mathbf{P}(\overline{\mathcal{A}(v_i, l(j)) \cap \mathcal{B}(v_i, l(j))}) \\
&\qquad\qquad + \mathbf{P}(\overline{\mathcal{A}(v_j, l(j)) \cap \mathcal{B}(v_j, l(j))})\Big)\Bigg) \\
\text{(by (7.5) and (7.6))} &= o(n^2).
\end{aligned} \tag{7.9}$$

Thus, (7.7) and (7.9) and Lemma 6.4 allow us to conclude that $X = \mathbf{E}[X](1 + o(1)) = n/2(1 + o(1))$ with probability tending to 1. □

The rest of the proof is devoted to establishing Lemmas 7.6 and 7.7. We first require a few geometric estimates.

Lemma 7.8. *Let $t \to 0$. If v is chosen uniformly in K, then $E(v,t) \subseteq K$ with probability $1 - O(t^2)$.*

Proof. We bound the failure probability. Note that $E(v,t) \not\subseteq K$ is equivalent to the event that the semimajor axis length of $E(v,t)$ exceeds half the distance of v to $\mathfrak{o}$ (the focal distance) by an amount greater than the distance from v_i to the boundary of K. This is established on the ball by noting that the curvature of the boundary is 1 everywhere, whereas the curvature at the extremes of the major axis of (γ-region) ellipses tends to infinity.

We label the semimajor and semiminor axis lengths a and b and let $r = \frac{\|v-\mathfrak{o}\|}{2}$ such that $a^2 = b^2 + r^2$.

Let us examine the excess $a-r$. Note that $(a^2-r^2) = (a+r)(a-r) \geq 2r(a-r)$, hence $a-r \leq \frac{a^2-r^2}{2r}$. Now the area of our ellipse is t, thus $\pi a\sqrt{a^2-r^2} = \pi ab = t$. An appeal to the quadratic formula yields

$$a^2 = \frac{1}{2}\left(r^2 + r^2\sqrt{1+\frac{4t^2}{\pi^2 r^4}}\right) \leq \left(r^2 + \frac{t^2}{\pi^2 r^2}\right). \tag{7.10}$$

Here we use the fact that, for all $x > 0$, we have $\sqrt{1+x} \leq 1 + x/2$. Thus we have $\frac{a^2-r^2}{2r} \leq \frac{t^2}{2\pi^2 r^3}$.

Fixing $r > r' > 0$ for some constant $r' < \frac{1}{2\sqrt{\pi}}$, we are thus assured that if v_i falls a distance at least $\frac{t^2}{2\pi^2(r')^3}(1+o(1))$ from the boundary of K then $E(v,t) \subseteq K$. For $r \leq r'$, $E(v,t)$ intersects the boundary of k only if $2a \geq 1/\sqrt{\pi}$, since the sum of the distances from any point in $\partial E(v,t)$ is $2a$. Now, as

$$a^2 = \frac{t^2}{\pi^2 a^2} + r^2 \leq \frac{t^2}{\pi^2 r^2} + r^2 \leq \frac{t^2}{\pi^2 (r')^2} + (r')^2,$$

we see that for t sufficiently small we have $a < 1/(2\sqrt{\pi})$.

Thus the only chance of failure comes from a point landing too close to the boundary of K, and, as above, this gives a failure probability of $O(t^2)$. $\square$

Proof of Lemma 7.6. For the moment, assume that v_i and v_j have the same length to the center $\mathfrak{o}$. Label the common focal length r (half the distance to the center), and label the farthest point of intersection of the two ellipses h. See Figure 4.

Consider the triangle on the right of Figure 4. Note that the sum of the lengths of the two smaller sides must equal the length of the major axis of the ellipse. Letting a be the semimajor axis length and x be the unlabeled edge in the figure, this implies that $2a = x + h$. Next, the law of cosines tells us that $h^2 + 4r^2 - 4rh\cos\frac{\theta}{2} = x^2$. Using our relation between x and h we obtain $h = \frac{a^2-r^2}{a - r\cos\frac{\theta}{2}}$.

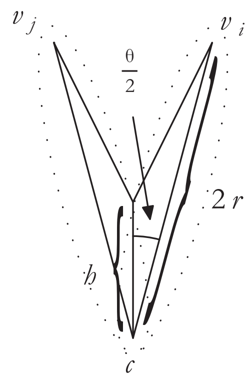

Figure 4. The intersection of two ellipses with a common focus and focal length forming an angle θ.

Recall that in Lemma 7.8 we already established the necessary asymptotics, so by (7.10) we have

$$a^2 = r^2\left(1 + \frac{t^2}{\pi^2 r^4}(1+o(1))\right),$$

under the assumption that $t^2/r^4 = o(1)$; i.e., $t = o(r^2)$. Taking square roots, we obtain the similar expansion $a = r(1 + \frac{t^2}{2\pi^2 r^4}(1+o(1)))$, again under the assumption $t = o(r^2)$. Thus we conclude that

$$h = \frac{a^2 - r^2}{a - r\cos\frac{\theta}{2}} \leq c\frac{t^2}{r_0^5(1-\cos\frac{\theta}{2})},$$

for some $c > 0$.

The case of v_i and v_j at different distances from $\mathfrak{o}$ follows by noting that lengthening one ellipse along its major axis while keeping the other fixed (maintaining the area in both) causes h to decrease. □

Now we need estimates on $\mathcal{A}(v_i), \mathcal{B}(v_i)$, etc.

Lemma 7.9. *For $l = \lfloor \lg\lg n \rfloor, \ldots, \lfloor \lg n \rfloor$ and $v_i, v_j \in V_n$, there exist constants $c_1, c_2, c_3, c_4 > 0$ such that*

$$\mathbf{P}(\overline{\mathcal{A}(v_i,l)}) \leq c_1 t_l^2,$$

$$\mathbf{P}(\overline{\mathcal{B}(v_i,l)}) \leq c_2 r_l^2 + (1 - t_l(1+o(1)))^{2^{l(i)}-1} + c_3 2^{l(i)} h_l^2,$$

$$\mathbf{P}(\overline{\mathcal{C}(v_i,v_j,l)}) = \theta_l.$$

In particular,

$$\mathbf{P}(\overline{\mathcal{D}(v_i, v_j, l)}) \leq 2c_1 t_l^2 + (1 - t_l(1 - o(1)))^{2^{l(i)}-1} + (1 - t_l(1 - o(1)))^{2^{l(j)}-1} + c_3 2^{l(i)} h_l^2 + c_3 2^{l(j)} h_l^2 + 2c_2 r_l^2 + c_4\theta_i.$$

Proof. The estimate on $\mathbf{P}(\overline{\mathcal{A}(v_i, l)})$ follows from Lemma 7.8.

The probability $\|v_i - \mathfrak{o}\| \leq r_l$ is πr_l^2. The probability that some point $v_{i'}$ such that $i' < i$ falls in $E(v_i, t_l) \cap B(v_i, h_l)$ or $E(v_i, t_l) \cap B(\mathfrak{o}, h_l)$ is at most $(i-1)2\pi h_l^2 \leq 2^{l(i)} 2\pi h_l^2$. Finally, the probability that every point $v_{i'}$ such that $i' < i$ misses $E(v_i, t_l)$ is at most $(1 - \text{Area}(E(v_i, t_l)))^{i-1} \leq (1 - t_l(1 - o(1)))^{2^{l(i)}-1}$. Here we use the fact that $2^{l(i)} - 1 \leq i$ and the bound $\text{Area}(E(v_i, t_l)) = t_l(1 - o(1))$ (the $o(1)$ term is due to those v_i that fall near the boundary of K). These facts combine to give the estimate on $\mathbf{P}(\overline{\mathcal{B}(v_i, l)})$.

The equation $\mathbf{P}(\overline{\mathcal{C}(v_i, v_j, l)})$ comes simply from the symmetry of K. □

We are now in a position to verify Lemma 7.7.

Proof of Lemma 7.7. We shall focus on Equation (7.5), the methods of which also give (7.4) and (7.6). Thus, we estimate $\sum_{i=1}^n \sum_{j>i}^n \mathbf{P}(\overline{\mathcal{D}(v_i, v_j, l(j))})$. As the RHS of all estimates in Lemma 7.9 depend only on relevent epochs, we can obtain the upper bound

$$\begin{aligned}
\sum_{i=1}^{n} \sum_{j>i}^{n} \mathbf{P}(\overline{\mathcal{D}(v_i, v_j, l(j))}) &\leq \sum_{i=\lfloor \lg\lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j=i+1}^{\lfloor \lg n \rfloor} 2^j\, \mathbf{P}(\overline{\mathcal{D}(v_{2^i}, v_{2^j})}) + O(n \lg n) \\
\text{(by Lemma 7.9)} \;&\leq \sum_{i=\lfloor \lg\lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j=i+1}^{\lfloor \lg n \rfloor} 2^j \Big(2c_1 t_j^2 + (1 - t_j(1 - o(1)))^{2^i - 1} \\
&\quad + (1 - t_j(1 - o(1)))^{2^j - 1} + c_3 2^i h_j^2 + c_3 2^j h_j^2 \\
&\quad + 2c_2 r_j^2 + c_4 \theta_j \Big) + O(n \lg n).
\end{aligned} \tag{7.11}$$

Here, the $O(n \lg n)$ term comes from neglecting those epochs for which Lemma 7.9 does not apply. As our goal is to bound the above by a $o(n^2)$ term we shall further neglect this error term and focus on the sum.

We evaluate the sum on each term. First, consider

$$\sum_{i=\lfloor \lg\lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j t_j^2 \leq \sum_{i=1}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j t_j^2 = \sum_{i=1}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j 2^{-j} = O(n \lg n).$$

Here, and in the remainder, it is understood that $i, j \leq \lfloor \lg n \rfloor$.

Similarly, we have

$$\sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j r_j^2 \leq \sum_{i=1}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j (1 - 2^{-1/7})^j = O(n),$$

as the inner sum is bounded by a geometric sum with modulus less than 1. The same estimate thus also shows $\sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j \theta_j = O(n)$.

Next, we compute

$$\sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j (2^j h_j^2) \leq \sum_{i=1}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} (2 - \varepsilon)^j = o(n^2),$$

given that the inner sum is bounded by $\sum_{j=1}^{\lfloor \lg n \rfloor} = O\left((2-\varepsilon)^{\lg n}\right) = o(n)$.

Finally, we check the sum

$$\sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j (1 - t_j(1 - o(1)))^{2^i - 1}.$$

We shall use the bound $t_j(1 - o(1)) \geq t_j/2$ for $j \geq \lfloor \lg \lg n \rfloor$. Thus,

$$\begin{aligned}
&\sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j (1 - t_j(1 - o(1)))^{2^i - 1} \\
&\qquad \leq \sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j \exp(-t_j(1 - o(1))(2^i - 1)) \\
&\qquad \leq \sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j \exp(-t_i(2^i - 1)/2) \\
&\qquad \leq O(n) \sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \exp(-2^{-i/2}(2^i - 1)/2) \\
&\qquad \leq O(n) \sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \exp(-2^{i/2-2}) = O(n),
\end{aligned}$$

by comparison to $\sum_{i=1}^{\infty} 2^i e^{2^{i/2-2}} < \infty$. Noting that $(1 - t_j(1 - o(1)))^{2^i - 1} \geq (1 - t_j(1 - o(1)))^{2^j - 1}$ as $j > i$, the estimate shows that $\sum_{i=\lfloor \lg \lg n \rfloor}^{\lfloor \lg n \rfloor} 2^i \sum_{j>i} 2^j (1 - t_j(1 - o(1)))^{2^j - 1}. = O(n)$ as well.

Thus, as (7.11) is a finite sum of sums each bounded by $o(n^2)$, (7.5) holds. The other estimates fall to the same methods. □

Proof of Theorem 4.5 (Part 2). With β fixed, we shall assume in what follows that $|\alpha - 1| = o(\gamma)$ for a parameter $\gamma \to 0$. Further, we assume that p is a point such that $\|\mathfrak{o} - p\| = r$ is bounded away from zero.

We consider the γ-region about p and construct inscribed and circumscribed figures about the γ-region. Let l_1 be the intersection of the line through $\mathfrak{o}$ and p and the γ region. The length of l_1 is thus $(1 + o(1))r$. Let l_2 be the segment perpendicular to l_1 that intersects the midpoint of $\overline{p\mathfrak{o}}$. The length of l_2 is $(1 + o(1))r\sqrt{\gamma}$. The convex hull of $l_1 \cup l_2$ forms a rhombus R_1 contained entirely in the γ-influence region with area $(1 + o(1))r^2\sqrt{\gamma}$. We can also form a circumscribed rectangle R_2 that is axis-parallel to segments l_1 and l_2 and has asymptotically twice the area of the rhombus.

Now let p be one of our randomly chosen points v_i, and further assume it has distance $\|v_i - \mathfrak{o}\|$ bounded away from 0 and $1/\sqrt{\pi}$, which happens with positive probability. Next, set $\gamma = (nr)^{-2}$, which causes the rhombus and rectangle to have areas $(1 + o(1))n^{-1}$ and $(1 + o(1))2n^{-1}$, respectively. Further, for n sufficiently great, the rhombus, ellipse, and rectangle all lie in K. Thus, the event that one of $\{v_1, \ldots, v_n\} - \{v_i\}$ falls in the rhombus and all others avoid the rectangle is at least

$$\binom{n}{1}\text{Area}(R_1)(1 - \text{Area}(R_2))^{n-2}$$
$$= \binom{n}{1}n^{-1}(1 + o(1))(1 - 2n^{-1}(1 + o(1)))^{n-2} > c_1 > 0,$$

for some positive constant c_1.

Now, for γ bounded above and hence the length of l_2 bounded above, we see that

$$\frac{\text{Area}(H(v_i, \beta) \cap R_1)}{\text{Area}(R_1)} > c_2 > 0.$$

Summarizing, the probability that there is exactly one point in the rhombus apart from v_i and that all other points fall outside the rectangle is bounded below by some positive constant. Further, the probability that a point chosen uniformly in the rhombus falls in $H(v_i, \beta)$ is bounded below by a positive constant. Thus, the total probability that the target of v_i is local is bounded below by a positive constant, which shows that $\mathbf{E}[X_i] > c > 0$ for some constant c, hence $\mathbf{E}[\rho(\beta)] > c > 0$. The upper bound is similar. □

Acknowledgments. The author wishes to thank Fan Chung for starting him on this project and for all her assistance. He would additionally like to thank Bill Chen and the Center for Combinatorics at Nankai University for their hospitality during part of this research. Finally, this work was made substantially more intelligible due to many thoughful comments of the anonymous referees.

References

[Alon and Spencer 00] Noga Alon and Joel Spencer. *The Probabilistic Method.* New York: John Wiley & Sons, Inc., 2000.

[Andersen et al. 05] R. Andersen, F. Chung, and L. Lu. "Modeling the Small-World Phenomena with Local Network Flow." *Internet Mathematics* 2:3 (2005), 359–385.

[Bai et al. 01] Z.-D. Bai, H.-K. Hwang, W.-Q. Liang, and T.-H. Tsai. "Limit Theorems for the Number of Maxima in Random Samples from Planar Regions." *Electronic Journal of Probability* 6 (2001), Paper no. 3. Available at http://www.math.washington.edu/~ejpecp/EjpVol6/paper3.pdf.

[Berger et al. 03] N. Berger, B. Bollobás, C. Borgs, J. Chayes, and O. Riordan. "Degree Distribution of the FKP Network Model." In *Automata, Languages and Programming: 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30–July 4, 2003, Proceedings*, Lecture Notes in Computer Science 2719, 725–738. Berlin: Springer, 2003.

[Bollobás and Chung 88] B. Bollobás and F. Chung. "The Diameter of a Cycle Plus a Random Matching." *SIAM Journal on Discrete Math* 1:3 (1988), 328–333.

[Carlson and Doyle 99] J. M. Carlson and J. Doyle. "Highly Optimized Tolerance: A Mechanism for Power Laws in Designed Systems." *Physical Review E* 60:2 (1999), 1412–1427.

[Carlson and Doyle 00] J. M. Carlson and J. Doyle. "Highly Optimized Tolerance: Robustness and Design in Complex Systems." *Physical Review Letters* 84:11 (2000), 2529–2532.

[Chung and Lu 02] F. Chung and L. Lu. "Connected Components in a Random Graph with Given Degree Sequences." *Annals of Combinatorics* 6, (2002), 125–145.

[Chung and Lu 06a] F. Chung and L. Lu. *Complex Graphs and Networks*, CBMS Regional Conference Series in Mathematics 107. Providence, RI: AMS, 2006.

[Chung and Lu 06b] F. Chung and L. Lu. "Concentration Inequalities and Martingale Inequalities, a Survey." *Internet Mathematics* 3:1 (2006), 79–127.

[Chung et al. 03] F. Chung, D. Jungreis, and S. Handjani. "Generalizations of Pólya's Urn Problem." *Annals of Combinatorics* 7:2 (2003), 141–153.

[Doyle et al. 05] John C. Doyle, David L. Alderson, Lun Li, Steven Low, Matthew Roughan, Stanislav Shalunov, Reiko Tanaka, and Walter Willinger. "The 'Robust Yet Fragile' Nature of the Internet." *Proceedings of the National Academy of Sciences* 102:41 (2005), 14497–14502. Available at (http://www.pnas.org/cgi/content/abstract/102/41/14497).

[Fabrikant et al. 02] A. Fabrikant, E. Koutsoupias, and C. Papadimitriou. "Heuristically Optimized Trade-offs: A New Paradigm for Powerlaws in the Internet." In *Automata, Languages and Programming: 29th International Colloquium, ICALP 2002, Mlaga, Spain, July 8-13, 2002, Proceedings* Lecture Notes in Computer Science 2380, 110–122. Berlin: Springer, 2002.

[Kleinberg 00] J. Kleinberg. J. Kleinberg. "The Small-World Phenomenon: An Algorithm Perspective." In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21–23, 2000, Portland, OR, USA*, pp. 163–170. New York: ACM Press, 2000.

[Liben-Nowell et al. 05] David Liben-Nowell, Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. "Geographic Routing in Social Networks." *Proceedings of the National Academy of Sciences* 102:33 (2005), 11623–11628.

[Smythe and Mahmoud 95] Robert T. Smythe and Hosam M. Mahmoud. "A Survey of Recursive Trees." *Theory of Probability and Mathematical Statistics* 51 (1995), 1–27. Translated from *Teoriya Imovirnostei ta Matematichna Statistika* 51 (1994), 1–29 (Ukrainian).

Ross M. Richardson, Department of Mathematics, University of California San Diego, La Jolla, CA 92093-0112, USA; and Center for Combinatorics, LPMC, Nankai University, Tianjin 300017, People's Republic of China (rmrichard@member.ams.org)

Received January 8 2007; accepted October 15, 2007.