Traveling in randomly embedded random graphs.

Alan Frieze*, Wesley Pegden†
Department of Mathematical Sciences,
Carnegie Mellon University,
Pittsburgh PA 15213,
U.S.A.

August 4, 2017

Abstract

We consider the problem of traveling among random points in Euclidean space, when only a random fraction of the pairs are joined by traversable connections. In particular, we show a threshold for a pair of points to be connected by a geodesic of length arbitrarily close to their Euclidean distance, and analyze the minimum length Traveling Salesperson Tour, extending the Beardwood-Halton-Hammersley theorem to this setting.

1 Introduction

The classical Beardwood-Halton-Hammersley theorem [2] (see also Steele [21] and Yukich [22]) concerns the minimum cost Traveling Salesperson Tour through n random points in Euclidean space. In particular, it guarantees the existence of an absolute (though still unknown) constant β_d such that if $X_1, X_2 \dots,$ is a random sequence of points, uniformly distributed in the d-dimensional cube $[0,1]^d$, the length $T(\mathcal{X}_{n,1})$ of a minimum tour through $X_1, \dots, X_n$ satisfies

$$T(\mathcal{X}_{n,1}) \approx \beta_d n^{\frac{d-1}{d}} \ a.s.$$

The present paper is concerned still with the problem of traveling among random points in Euclidean space. In our case, however, we suppose that only a (random) subset of the pairs of points are joined by traversable connections, independent of the geometry of the point set.

In particular, we study random embeddings of the Erdős-Rényi-Gilbert random graph $G_{n,p}$ into the d-dimensional cube $[0,1]^d$. We let $\mathcal{X}_n$ denote a uniformly random set of points $X_1, X_2, \dots, X_n \in [0,1]^d$, and we denote by $\mathcal{X}_{n,p}$ the random graph whose vertex set is $\mathcal{X}_n$ and whose pairs of vertices are joined by edges each with independent probability p. Edges are weighted by the Euclidean distance between their points, and we are interested in the total edge-weight required to travel about the graph.

*Research supported in part by NSF grant DMS1362785
†Research supported in part by NSF grant DMS

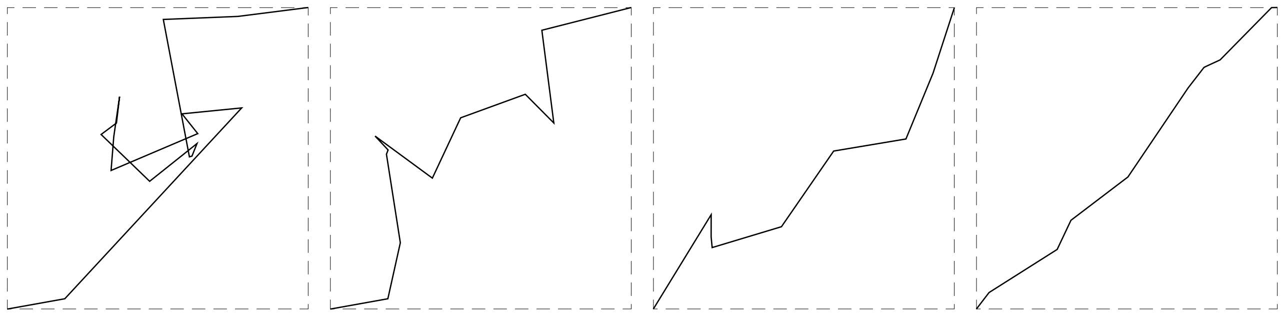

Figure 1: Paths in an instance of $\mathcal{X}_{n,p}$ for $d=2$, $n=2^{30}$, and $p=\frac{10}{n}, \frac{25}{n}, \frac{50}{n}$, and $\frac{200}{n}$, respectively. In each case, the path drawn is the shortest route between the vertices x and y which are closest to the SW and NE corners of the square. (See Q. 2, Section 5.)

This model has received much less attention than the standard model of a random geometric graph, defined as the intersection graph of unit balls with random centers $X_i, i \in [n]$, see Penrose [17]. We are only aware of the papers by Mehrabian [14] and Mehrabian and Wormald [15] who studied the *stretch factor* of $\mathcal{X}_{n,p}$. In particular, let $||x-y||$ denote the Euclidean distance between vertices x, y, and $\text{dist}(x,y)$ denote their distance in $\mathcal{X}_{n,p}$. They showed (considering the case $d=2$) that if $n(1-p) \to \infty$, then the stretch factor

$$\sup_{x,y \in \mathcal{X}_{n,p}} \frac{\text{dist}(x,y)}{||x-y||}$$

tends to ∞ with n.

As a counterpoint to this, our first result shows a very different phenomenon when we pay attention to additive rather than multiplicative errors. In particular, for $p \gg \frac{\log^d n}{n}$, the distance between a typical pair of vertices is arbitrarily close to their Euclidean distance, while for $p \ll \frac{\log^d n}{n(\log\log n)^{2d}}$, the distance between a typical pair of vertices in $\mathcal{X}_n$ is arbitrarily large (Figure 1). (We write $\log^k x$ for $(\log x)^k$.) In particular, this means that when $\frac{\log^d n}{n} \ll p < 1-\varepsilon$, the supremum in the stretch factor theorem of Mehrabian and Wormald is due just to pairs of vertices which are very close together.

Theorem 1.1. *Let $\omega = \omega(n) \to \infty$. We have for $d \geq 2$:*

(a) For $p \leq \frac{1}{\omega^d (\log\log n)^{2d}} \frac{\log^d n}{n}$ and $u = X_1$, $v = X_2$, we have

$$\text{dist}(u,v) \geq \frac{\omega}{8de^d}, \qquad a.a.s.^1$$

(b) For $p \geq \frac{\omega \log^d n}{n}$, we have a.a.s. that uniformly for all pairs of vertices $u, v \in \mathcal{X}_n$,

$$\text{dist}(u,v) = ||u-v|| + o(1).$$

Theorem 1.1 means that, even for p quite small, it is not that much more expensive to travel from one vertex of $\mathcal{X}_{n,p}$ to another than it is to travel directly between them in the plane. On the other hand, there is a dramatic dependence on p if the goal is to travel among *all* points. Let $T(\mathcal{X}_{n,p})$ denote the length of a minimum length Traveling Salesperson tour in $\mathcal{X}_{n,p}$, i.e. a minimum length Hamilton cycle.

[1] A sequence of events $\mathcal{E}_n$ occurs *asymptotically almost surely* (a.a.s.) if $\lim_{n\to\infty} \mathbf{Pr}(\mathcal{E}_n) = 1$.

Theorem 1.2. *There exists a sufficiently large constant $K > 0$ such that for all $p = p(n)$ such that $p \geq \frac{K \log n}{n}$, $d \geq 2$, we have that*

$$T(\mathcal{X}_{n,p}) = \Theta\left(\frac{n^{\frac{d-1}{d}}}{p^{1/d}}\right) \qquad a.a.s. \tag{1}$$

(Recall that $f(n) = \Theta(g(n))$ means that $f(n)$ is bounded between positive constant multiples of $g(n)$ for sufficiently large n.) As the threshold for $G_{n,p}$ to be Hamiltonian is at $p = \frac{\log n + \log\log n + \omega(n)}{n}$ (see e.g. Bollobás [3]), this theorem covers nearly the entire range of p for which a TSP tour exists a.a.s.

Finally, we extend the asymptotically tight BHH theorem [2] to the case of $\mathcal{X}_{n,p}$ for any constant p. To formulate an "almost surely" statement, we let $\mathcal{X}_{\mathcal{N},p}$ denote a random graph on a random embedding of $\mathcal{N} = \{1, 2, \ldots, \}$ into $[0,1]^d$, where each pair $\{i, j\}$ is independently present as an edge with probability p, and consider $\mathcal{X}_{n,p}$ as the restriction of $\mathcal{X}_{\mathcal{N},p}$ to the first n vertices $\{1, \ldots, n\}$.

Theorem 1.3. *If $d \geq 2$ and $p > 0$ is constant, then there exists $\beta_{d,p} > 0$ such that*

$$T(\mathcal{X}_{n,p}) \approx \beta_{d,p} n^{\frac{d-1}{d}} \qquad a.s.$$

Karp's algorithm [13] for a finding an approximate tour through $\mathcal{X}_n$ extends to the case $\mathcal{X}_{n,p}$, p constant as well:

Theorem 1.4. *For fixed $d \geq 2$ and p constant, then there is an algorithm that a.s. finds a tour in $\mathcal{X}_{n,p}$ of value $(1 + o(1))\beta_{d,p} n^{(d-1)/d}$ in polynomial time, for all $n \in \mathcal{N}$.*

2 Traveling between pairs

In this section, we prove Theorem 1.1.

2.1 Proof of Theorem 1.1(a)

Outline of proof

This is straightforward. We show by the first moment method that any path between u and v with "many" edges must contain a significant number of "long" edges and hence must be as long as claimed. We then show that a.a.s. there are no paths between u and v without many edges.

Proof proper

Let ν_d denote the volume of a d-dimensional unit ball; recall that ν_d is bounded ($\nu_d \leq \nu_5 < 6$ for all d).

Let an edge be *long* if its length is at least $\ell_1 = \frac{\omega(\log\log n)^2}{4e^d \log n}$. Let $\varepsilon = \frac{1}{\log\log n}$ and let $\mathcal{A}$ be the event that there exists a path with k edges, $k \geq k_0 = \frac{\log n}{2d \log\log n}$ from u to v that uses at most εk long edges. Then

$$\text{(2)}\qquad \mathbf{Pr}(\mathcal{A}) \leq \sum_{k\geq k_0}(k-1)!\binom{n}{k-1}p^k\binom{k}{\varepsilon k}\left(\nu_d\left(\frac{\omega(\log\log n)^2}{4e^d\log n}\right)^d\right)^{(1-\varepsilon)k}$$
$$\leq \sum_{k\geq k_0} n^{k-1}p^k\left(\frac{e}{\varepsilon}\right)^{\varepsilon k}\left(\nu_d\left(\frac{\omega(\log\log n)^2}{4e^d\log n}\right)^d\right)^{(1-\varepsilon)k}$$
$$\leq \frac{1}{n}\sum_{k\geq k_0}\left(\frac{\nu_d^{1-\varepsilon}\log^{d\varepsilon} n}{(4e^d)^{d(1-\varepsilon)}}\cdot\left(\frac{e}{\varepsilon}\right)^{\varepsilon}\right)^k$$
$$\leq \frac{1}{n}\sum_{k\geq k_0}\left(\frac{6e^{d+o(1)}}{4e^{d^2}}\right)^k = o(1),$$

after using $d \geq 2$ and $\log^{\varepsilon} n = e$.

Explanation of (2)**:** Choose the $k-1$ interior vertices of the possible path and order them in one of $(k-1)!\binom{n}{k-1}$ ways as $(u_1, u_2, \ldots, u_{k-1})$. Then p^k is the probability that the edges exist in $G_{n,p}$. Now choose the short edges $e_i = (u_{i-1}, u_i), i \in I$ in one of $\binom{k}{(1-\varepsilon)k} = \binom{k}{\varepsilon k}$ ways and bound the probability that these edges are short by $\left(\nu_d\left(\frac{\omega(\log\log n)^2}{4e^d\log n}\right)^d\right)^{(1-\varepsilon)k}$ viz. the probability that u_i is mapped to the ball of radius ℓ_1, center u_{i-1} for $i \in I$.

Now a.a.s. the shortest path in $G_{n,p}$ from u to v requires at least k_0 edges: Indeed the expected number of paths of length at most k_0 from u to v can be bounded by

$$\sum_{k=1}^{k_0}(k-1)!\binom{n}{k-1}p^k \leq \frac{1}{n}\sum_{k=1}^{k_0}\left(\frac{\log^d n}{\omega^d(\log\log n)^{2d}}\right)^k = o(1).$$

So a.a.s.

$$\mathrm{dist}(u,v) \geq \varepsilon k_0\ell_1 = \frac{\varepsilon\log n}{2d\log\log n}\cdot\frac{\omega(\log\log n)^2}{4e^d\log n} = \frac{\omega}{8de^d}.$$

□

2.2 Proof of Theorem 1.1(b)

Outline of proof

We first consider two points u, v such that $||u-v|| \geq \gamma = \frac{1}{\log\log n}$. We then consider a set of small disjoint balls with centers on the line joining u, v. We argue that a.a.s. (i) all of these balls contain (relatively) giant components, (ii) there is an edge joining any pair of giant components inside each ball, (iii) the diameter of each of these giant components is small and (iv) there is an edge between u and one of the g giant components X closest to u and an edge between v and one of the g giant components Y closest to v. This gives a path consisting of an edge from u to the giant component X plus a walk inside X plus an edge to the giant component Y plus an edge to v. Because the balls are small the length of this path is close to $||u-v||$. We reduce the case where $||u-v|| \leq \gamma$ to the first case.

Proof proper

We begin by considering the case of vertices u, v at distance $||u-v|| \geq \gamma$. Letting $\delta = \frac{1}{\log n}$, then, for

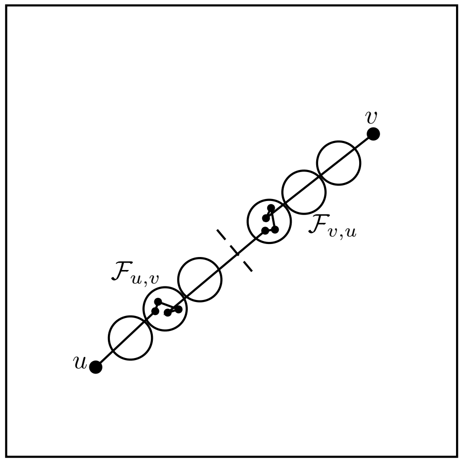

Figure 2: Finding a short path.

sufficiently large n, we can find a set $\mathcal{B}$ of at least $\frac{2C}{\delta}$, $C = \frac{\gamma}{8}$, disjoint balls of radius δ centered on the line from u to v, such that $\frac{C}{\delta}$ of the balls are closer to u than v, and $\frac{C}{\delta}$ balls are closer to v than u (Figure 2). Denote these two families of $\frac{C}{\delta}$ balls by $\mathcal{F}_{u,v}$ and $\mathcal{F}_{v,u}$. (The sets $\mathcal{B}$, $\mathcal{F}_{u,v}$ and $\mathcal{F}_{v,u}$ are fixed for the rest of the argument.)

Given a ball $B \in \mathcal{F}_{\{u,v\}} = \mathcal{F}_{u,v} \cup \mathcal{F}_{v,u}$, the induced subgraph G_B on vertices of $\mathcal{X}$ lying in B is a copy of $G_{N,p}$, where $N = N(B)$ is the (random) number of vertices lying in B. Let

$$\mathcal{S}_B \text{ be the event that } N(B) \in \left[\frac{N_0}{2^{d+1}}, 2N_0\right] \text{ where } N_0 = \nu_d \delta^d n.$$

(Dividing by 2^{d+1} accounts for points close to the boundary of $[0,1]^d$.)

Now $N(B)$ is distributed as the binomial $Bin(n,q)$ where $q \in \nu_d \delta^d [2^{-d}, 1]$. The following Chernoff bounds will thus be useful:

(3) $$\mathbf{Pr}(Bin(M,p) \leq (1-\varepsilon)Mp) \leq e^{-\varepsilon^2 Mp/2} \text{ for } 0 \leq \varepsilon \leq 1.$$

(4) $$\mathbf{Pr}(Bin(M,p) \geq (1+\varepsilon)Mp) \leq e^{-\varepsilon^2 Mp/3} \text{ for } 0 \leq \varepsilon \leq 1.$$

The bounds (3) and (4) imply that for $B \in \mathcal{F}_{\{u,v\}}$,

$$\mathbf{Pr}\left(\neg \mathcal{S}_B\right) \leq e^{-\Omega(n\delta^d)} = e^{-n^{1-o(1)}}.$$

This gives us that a.a.s. $\mathcal{S}_B$ occurs for all pairs $u, v \in \mathcal{X}$ with $||u - v|| \geq \gamma$. We now argue that for all $B \in \mathcal{B}$:

(A) All subgraphs G_B for $B \in \mathcal{F}_{\{u,v\}}$ have a giant component X_B, containing at least $N_0/2^{d+2}$ vertices. Indeed, the expected average degree in G_B is $Np = \Omega(\omega) \to \infty$ (and with probability $1 - e^{-n^{1-o(1)}}$ we have $N = n^{1-o(1)}$) and at this value the giant component is almost all of B a.a.s. In particular, since $\mathcal{S}_B$ occurs, we have that

(5) $$\mathbf{Pr}(|X_B| \leq N_0/2^{d+2} \mid \mathcal{S}_B) \leq e^{-\Omega(N_0)} \leq e^{-\Omega(\delta^d n)} = o(n^{-3}).$$

See [3] for the first inequality in (5). This can be inflated by $n^2 \cdot (2C\log n)$ to account for pairs u, v and the choice of $B \in \mathcal{F}_{\{u,v\}}$.

(B) There is an edge between X_B and $X_{B'}$ for all $B, B' \in \mathcal{F}_{\{u,v\}}$.
Indeed, the probability that there is no edge between $X_B, X_{B'}$, given (A), is at most

$$(1-p)^{N_0^2/2^{2d+2}} \leq e^{-\Omega(\delta^{2d}n^2p)} \leq e^{-n^{1-o(1)}}.$$

This can be inflated by $n^2 \cdot (C\log n)^2$ to account for all pairs u, v and all pairs B, B'.

(C) For each $B \in \mathcal{F}_{\{u,v\}}$, the graph diameter $\mathrm{diam}(X_B)$ (the maximum number of edges in any shortest path in X_B) satisfies

$$\mathbf{Pr}\left(\mathrm{diam}(X_B) > \frac{100\log N_0}{\log(N_0p)}\right) \leq n^{-3}.$$

This can be inflated by $n^2 \cdot (2C\log n)$ to account for pairs u, v and the choice of $B \in \mathcal{F}_{\{u,v\}}$. Fernholz and Ramachandran [5] and Riordan and Wormald [20] gave tight estimates for the diameter of the giant component, but we need this cruder estimate with a lower probability of being exceeded. We prove this later in Lemma 2.1. It will be convenient for the proof of Lemma 2.1 to assume that $N_0p = O(\log N_0)$. There is no loss in generality because Theorem 1.1(b) holds a fortiori for larger p. This follows from a standard coupling argument, involving adding random edges to increase the edge probability.

Part (C) implies that with high probability, for any u, v at distance $\geq \gamma$ and all $B \in \mathcal{F}_{\{u,v\}}$ and vertices $x, y \in X_B$,

$$\text{(6)} \qquad \mathrm{dist}(x,y) \leq 200\delta \times \frac{\log N_0}{\log(N_0p)} \leq \frac{200}{\log n} \times \frac{\log n - d\log\log n + \log \nu_d}{\log \omega + \log \nu_d} = o(1).$$

As the giant components X_B ($B \in \mathcal{F}_{u,v}$) contain in total at least $\frac{C}{\delta}\frac{N_0}{2^{d+2}} = \frac{C}{2^{d+2}}\nu_d n\delta^{d-1}$ vertices, the probability that u has no neighbor in these giant components is at most

$$(1-p)^{C\nu_d n\delta^{d-1}/2^{d+2}} \leq e^{-C\nu_d np\delta^{d-1}/2^{d+2}} = n^{-\omega C\nu_d/2^{d+2}}.$$

In particular, the probability is small after multiplication by n^2, and thus a.a.s., for all pairs $x, y \in X_{n,p}$, x has a neighbor in X_B for some $B \in \mathcal{F}_{u,v}$ and y has a neighbor in $X_{B'}$ for some $B' \in \mathcal{F}_{v,u}$. Now by part (B) and equation (6), we can find a path

$$u, w_0, w_1, \ldots, w_s, z_t, z_{t-1}, \ldots, z_1, z_0, v$$

from u to v where the w_i's are all in some X_B for $B \in \mathcal{F}_{u,v}$ and the total Euclidean length of the path $w_0, \ldots, w_s$ tends to zero with n, and the z_i's are all in some $X_{B'}$ for some $B' \in \mathcal{F}_{v,u}$, and the total Euclidean length of the path $z_0, \ldots, w_t$ tends to zero with n. Meanwhile, the Euclidean segments corresponding to the three edges u, w_0, w_s, z_t, and z_0, v lie within δ of disjoint segments of the line segment from u to v, and thus have total length $\leq ||u-v|| + 6\delta$, giving

$$\text{(7)} \qquad \mathrm{dist}(u,v) \leq ||u-v|| + 6\delta + o(1) = ||u-v|| + o(1).$$

We must also handle vertices $u, v \in \mathcal{X}_{n,p}$ with $||u-v|| < \gamma$. Given such a pair, we let B_u, B_v denote any choice of balls of radius γ such $\mathrm{dist}(B_u, B_v) \geq \gamma$, $\mathrm{dist}(B_u, u), \mathrm{dist}(B_v, v) \leq \gamma(\sqrt{d}+2)$. (These bounds are chosen to make such a choice trivially possible, even when u, v are close to a corner.) Observe that we have: where C_u, C_v denote the giant components of B_u, B_v, and $x \sim y$ means that $\{x, y\}$ is an edge of $\mathcal{X}_{n,p}$,

$$\mathbf{Pr}(\forall u, v \in \mathcal{X}_{n,p}, \exists w \in C_u, z \in C_v \text{ such that } u \sim w, v \sim z) \to 1$$

with n since a.a.s we have that B_u and B_v contain at least $\nu_d n\gamma^d/2^{d+2}$ points for all $u, v \in \mathcal{X}_{n,p}$ and we have that $1 - 2n^2(1-p)^{n\cdot\nu_d\gamma^d/2^{d+2}} \to 1$. In particular, we can a.a.s for all pairs $u, v \in \mathcal{X}_{n,p}$ find $w \sim u$ within distance $\gamma(\sqrt{d}+4)$ of u, $z \sim v$ within Euclidean distance $\gamma(\sqrt{d}+4)$ of v, such that

$$\gamma \leq ||w - z|| \leq (2\sqrt{d}+8)\gamma.$$

Now, we can use the previous case (7) to see that

$$\text{dist}(u, v) \leq (2\sqrt{d}+9)\gamma + 6\delta + o(1) = o(1).$$

In particular, $\text{dist}(u, v) - ||u - v|| = o(1)$. □

We complete the proof of Theorem 1.1 by proving

Lemma 2.1. *Suppose that $Np = \omega \to \infty, \omega = O(\log N)$ and let C_1 denote the unique giant component of size $N - o(N)$ in $G_{N,p}$, that q.s.*[2] *exists. Then for L large,*

$$\mathbf{Pr}\left(\text{diam}(C_1) \geq \frac{L\log N}{\log Np}\right) \leq O(N^{-L/10}).$$

Proof. Let $\mathcal{B}(k)$ be the event that there exists a set S of k vertices in $G_{N,p}$ that induces a connected subgraph in which more than half of the vertices have less than $\omega/2$ neighbors outside S. Then for $k = o(N)$ we have

$$\begin{aligned}
\mathbf{Pr}(\mathcal{B}(k)) &\leq \binom{N}{k} p^{k-1} k^{k-2} 2^k\, \mathbf{Pr}(Bin(N-k,p) \leq \omega/2)^{k/2} && (8)\\
&\leq \frac{e^k\omega^k}{pk^2} 2^k \left(e^{-((N-k)p-\omega/2)^2/(2(N-k)p)}\right)^{k/2}, \text{ from (3) with } \varepsilon = 1 - \frac{\omega}{2(N-k)p},\\
&\leq \frac{e^k\omega^k}{pk^2} 2^k \left(e^{-(.99\omega-\omega/2)^2/2\omega}\right)^{k/2}\\
&\leq p^{-1}(2e\omega e^{-\omega/20})^k \leq Ne^{-k\omega/21}. && (9)
\end{aligned}$$

Explanation of (8)**:** $\binom{N}{k}$ bounds the number of choices for S. We then choose a spanning tree T for S in k^{k-2} ways. We multiply by p^{k-1}, the probability that T exists. We then choose half the vertices X of S in at most 2^k ways and then multiply by the probability that each $x \in X$ has at most $\omega/2$ neighbors in $[N] \setminus S$.

If $\kappa = \kappa(L) = \frac{L\log N}{\log Np}$ then (9) implies that $\mathbf{Pr}(\mathcal{B}(\kappa)) \leq N^{1-L}$.

[2] A sequence of events $\mathcal{E}_n$ occurs *quite surely* q.s. if $\mathbf{Pr}(\neg\mathcal{E}_n) = O(n^{-K})$ for all positive constants K.

Next let $\mathcal{D}(k) = \mathcal{D}_N(k)$ be the event that there exists a set S of size k for which the number of edges $e(S)$ contained in S satisfies $e(S) \geq 2k$. Then,

$$\mathbf{Pr}(\mathcal{D}(k)) \leq \binom{N}{k}\binom{\binom{k}{2}}{2k}p^{2k} \leq \left(\frac{Ne}{k} \cdot \left(\frac{ke\omega}{2N}\right)^2\right)^k = \left(\frac{ke^3\omega^2}{4N}\right)^k.$$

Since $\omega = O(\log N)$ we have that q.s.

$$\nexists k \in [\kappa(1), N^{3/4}] \text{ such that } \mathcal{D}(k) \text{ occurs.}$$

Now let $\mathcal{B}(k_1, k_2) = \bigcup_{k=k_1}^{k_2} \mathcal{B}(k)$ and $\mathcal{D}(k_1, k_2) = \bigcup_{k=k_1}^{k_2} \mathcal{D}(k)$, and suppose that

$$\mathcal{B}(k_1, k_2) \cup \mathcal{D}(k_1, k_2) \text{ does not occur,}$$

where $k_1 = \kappa(L/4)$ and $k_2 = N^{3/4}$. Fix a pair of vertices v, w and define sets $S_0, S_1, S_2, \dots$ where S_i is the set of vertices at distance i from v. If there is no $i \leq k_1$ with $w \in S_i$ then we must have $S_{k_1} \neq \emptyset$ and $|S_{\leq k_1}| \geq k_1$ where $S_{\leq t} = \bigcup_{i=0}^{t} S_i$ for $t \geq 0$. This is because $v, w \in C_1$ and C_1 is connected and so $|S_{\leq i+1}| \geq |S_{\leq i}| + 1$. We also see that $k_1 \leq |S_{\leq t}| \leq N^{3/4}$ implies that $|S_{t+1}| \geq \omega |S_{\leq t}|/10$. Indeed, if $|S_{t+1}| < \omega |S_{\leq t}|/10$ then $S_{\leq t+1}$ has at most $(\omega + 10)|S_{\leq t}|/10$ vertices and more than $\omega |S_{\leq t}|/4$ edges, contradiction.

Thus if L is large, then we find that there exists $t \leq k_1 + \kappa(3/4) \leq N^{3/4}$ such that $|S_{\leq t}| \geq N^{3/4}$ and so also that $|S_t| \geq (1 - o(1))N^{3/4}$. Now apply the same argument from w to create sets $T_0, T_1, \dots, T_s$, where either we reach v or find that $|T_s| \geq N^{3/4}$ where $s \leq k_1 + \kappa(3/4)$. At this point the edges between S_t and T_s are unconditioned and the probability there is no $S_t : T_s$ edge is at most $(1-p)^{N^{3/2}} = O(e^{-\Omega(N^{1/2})})$. □

3 Traveling among all vertices

Our first aim is to prove Theorem 1.3; this will be accomplished in Section 3.2, below. In fact, we will prove the following general statement, which will also be useful in the proof of Theorem 1.2:

Theorem 3.1. *Let $\mathcal{Y}_1^d \subset [0,1]^d$ denote a set of points chosen from any fixed distribution, such that the cardinality $Y = |\mathcal{Y}_1^d|$ satisfies $\mathbf{E}(Y) = \mu > 0$ and $\mathbf{Pr}(Y \geq k) \leq C\rho^k$ for all k, for some $C > 0, \rho < 1$. For $t > 0$ let $\mathcal{Y}_t^d$ denote a random set of points in $[0,t]^d$ obtained from the union of t^d independent copies $\mathcal{Y}_1^d + x$ ($x \in \{0, \cdots, t-1\}^d$).*

If $p > 0$ is constant, $d \geq 2$, and $\mathcal{Y}_{t,p}^d$ denotes the random graph on $\mathcal{Y}_t^d$ with independent edge probabilities p, then $\exists \beta > 0$ (depending on p and the process generating $\mathcal{Y}_1^d$) such that

(i) $T(\mathcal{Y}_{t,p}^d) \approx \beta t^d$ *a.a.s., and*

(ii) $T(\mathcal{Y}_{t,p}^d) \leq \beta t^d + o(t^d)$ *q.s.*[3]

[3] In this context $O(n^{-\omega(1)})$ is replaced by $O(t^{-\omega(1)})$.

Note that as a probabilistic statement, Part (i) above asserts that there exists a choice for $o(1)$ (a function of t, say, tending to 0) such that $(1-o(1))\beta t^d \leq T(\mathcal{Y}^d_{t,p}) \leq (1+o(1))\beta t^d$ holds a.a.s. Similarly for Part (ii), the statement asserts the existence of a suitable fixed choice of $o(t^d)$ (a function of t, whose ratio to t^d tends to 0).

The restriction $\mathbf{Pr}\left(|\mathcal{Y}^d_1| \geq k\right) \leq C\rho^k$ simply ensures that we have exponential tail bounds on the number of points in a large number of independent copies of $\mathcal{Y}^d_1$:

Observation 3.2. *For the total number T_n of points in n independent copies of $\mathcal{Y}^d_1$, we have for some absolute constant $A_{C,\rho} > 0$,*

$$\mathbf{Pr}(|T_n - \mu n| > \delta\mu n) < e^{-A_{C,\rho}\delta^2\mu^2 n}.$$

□

This is straightforward to prove, but we do not have a reference and so we give a sketch proof in the appendix.

Note that the conditions on the distribution of $\mathcal{Y}^d_t$ are satisfied for a Poisson cloud of intensity 1, and it is via this case that we will derive Theorem 1.3. Other examples for which these conditions hold include the case where $\mathcal{Y}^d_t$ is simply a suitable grid of points, or is a random subset of a suitable grid of points in $[0,t]^d$, and we will make use of this latter case of Theorem 3.1 in our proof of Theorem 1.2.

Outline of proof of Theorem 3.1
Our proof uses subadditivity, but some of the standard properties of the classical case (e.g., monotonicity) fail in our setting, requiring us to use induction on d to achieve the result. For technical reasons (see also Question 4 of Section 5) Theorems 3.1 and 1.3 are given just for $d \geq 2$, and before beginning with the induction, we must carry out a separate argument to bound the length of the tour in 1 dimension.

When $d = 1$ all we can prove is an $O(n)$ bound on the length of the minimum tour. We do this by examining a natural greedy algorithm for finding a tour. This is the content of Lemma 3.4. After this we prove a sort of Lipschitz condition for the tour length, see Lemma 3.7. This will substitute for monotonicity. After this we can push ahead using subadditivity.

3.1 Bounding the expected tour length in 1 dimension

We begin with the following simple lemma.

Lemma 3.3. *Let σ be a permutation of $[n]$, and let $\ell(\sigma) = \sum_{i=1}^{n-1} |\sigma_{i+1} - \sigma_i|$. Then*

$$\ell(\sigma) < \sigma_n + 4 \cdot \mathrm{inv}(\sigma),$$

where $\mathrm{inv}(\sigma)$ *is the number of inversions in σ (i.e. $|\{(i,j) : i < j$ and $\sigma_i > \sigma_j\}|$).*

Proof. We prove this by induction on n. It is trivially true for $n = 1$ since in this case $\ell(\sigma) = 0$. Assume now that $n > 1$, and given a permutation σ of $[n]$, consider permutation σ' of $[n-1]$ obtained by truncation:

$$\sigma'_i = \begin{cases} \sigma_i & \text{if } \sigma_i < \sigma_n \\ \sigma_i - 1 & \text{if } \sigma_i > \sigma_n \end{cases}$$

We have by induction that

$$\ell(\sigma') \leq \sigma'_{n-1} + 4 \cdot \text{inv}(\sigma'). \tag{10}$$

Now observe that

$$\begin{aligned}\ell(\sigma) &= \ell(\sigma') + |\sigma_n - \sigma_{n-1}| + |\{i|\sigma_i < \sigma_n < \sigma_{i+1} \text{ OR } \sigma_i > \sigma_n > \sigma_{i+1}\}| \\ &\leq \ell(\sigma') + |\sigma_n - \sigma_{n-1}| + 2(\text{inv}(\sigma) - \text{inv}(\sigma')),\end{aligned}$$

since $|\{i|\sigma_i < \sigma_n < \sigma_{i+1}\}|$ and $|\{i|\sigma_{i+1} < \sigma_n < \sigma_i\}|$ are each bounded by $\text{inv}(\sigma) - \text{inv}(\sigma')$. Recalling that $\text{inv}(\sigma) = \text{inv}(\sigma^{-1})$, we have

$$\text{inv}(\sigma) - \text{inv}(\sigma') = n - \sigma_n.$$

Since $\sigma'_{n-1} \leq \sigma_{n-1}$, (10) gives that

$$\begin{aligned}\ell(\sigma) &\leq \sigma_{n-1} + 4 \cdot \text{inv}(\sigma') + |\sigma_n - \sigma_{n-1}| + 2(\text{inv}(\sigma) - \text{inv}(\sigma')) \\ &= \sigma_{n-1} + 2 \cdot \text{inv}(\sigma') + 2(\text{inv}(\sigma) - n + \sigma_n) + |\sigma_n - \sigma_{n-1}| + 2(\text{inv}(\sigma) - \text{inv}(\sigma')) \\ &= \sigma_{n-1} + 4 \cdot \text{inv}(\sigma) - 2n + 2\sigma_n + |\sigma_n - \sigma_{n-1}| \\ &= \sigma_n + 4 \cdot \text{inv}(\sigma) - (2n - \sigma_n - \sigma_{n-1} - |\sigma_n - \sigma_{n-1}|) \\ &\leq \sigma_n + 4 \cdot \text{inv}(\sigma).\end{aligned}$$

□

For the 1-dimension case of Theorem 1.3, we have, roughly speaking, a 1-dimensional string of points joined by some random edges. Lemma 3.3 allows us to prove the following lemma, which begins to approximate this situation.

Lemma 3.4. *Consider the random graph $G = G_{n,p}$ on the vertex set $[n]$ with constant p, where each edge $\{i,j\} \in E(G)$ is given length $|i-j| \in \mathbb{N}$. Let Z denote the minimum length of a Hamilton cycle in G starting at vertex 1, assuming one exists. If no such cycle exists let $Z = n^2$. Then there exists a constant A_p such that*

$$\mathbf{E}(Z) \leq A_p n \text{ and } Z \leq A_p n, \text{ q.s.}$$

Proof. We first write $G = G_1 \cup G_2 \cup G_3$ where the G_i are independent copies of G_{n,p_1}, where $1 - p = (1-p_1)^3$. Note that $p_1 \geq p/3$. We consider G_1 to be the restriction to $[n]$ of a random graph $\mathcal{G} = G_{\mathcal{N},p_1}$ on the set of all natural numbers, and will begin by constructing a path in $\mathcal{G}$ via the following algorithm: We start with $v_1 = 1$. Then for $j \geq 1$ we let

$$\phi(j) = \min\{k \in \mathcal{N} : k \notin \{v_1, v_2, \ldots, v_j\} \text{ and } \{v_j, k\} \in E(G_1)\}$$

and let $v_{j+1} = \phi(j)$ i.e. we move from v_j to the lowest index k that has not been previously visited. This constructs an infinite path in $\mathcal{N}$, and we define j_0 by

$$j_0 = \max\{j \in \mathcal{N} : i \leq j \implies v_i \leq n\}.$$

In particular, observe that $P_1 = v_1, v_2, \ldots, v_{j_0}$ is a path in G_1 of length $\Lambda_1 = \sum_{j=1}^{j_0-1} |v_{j+1} - v_j|$. It is convenient to extend the sequence $v_1, \ldots, v_{j_0}$ to a permutation of $[n]$; to do this, we let $\sigma_i = v_i$ for $i \leq j_0$, and then let $\sigma_{j_0+1}, \ldots, \sigma_n$ be $[n] \setminus \{v_1, \ldots, v_{j_0}\}$ in increasing order. Applying Lemma 3.3 to σ gives that the length Λ_1 of the initial part corresponding to the path is at most $\ell(\sigma) < n + 4 \cdot \text{inv}(\sigma)$. So we would like to bound $\text{inv}(\sigma)$.

Observe first that $\mathbf{Pr}(j_0 \le n-k) \le n(1-p_1)^k$. This is because at j_0 we find that v_{j_0} has no neighbors in the set of unvisited vertices and the existence of such edges is unconditioned at this point. So,

$$j_0 \ge n - \frac{\log^2 n}{p_1} \; q.s. \tag{11}$$

Now let $\alpha_j = |\{i > j : \sigma_i < \sigma_j\}|$ for all $1 \le j \le n$, so that $\mathrm{inv}(\sigma) = \alpha_1 + \alpha_2 + \cdots + \alpha_n$. In fact, by our definition of σ, we have that $\alpha_j = 0$ for all $j > j_0$, giving

$$\mathrm{inv}(\sigma) = \alpha_1 + \alpha_2 + \cdots + \alpha_{j_0},$$

since σ contains no inversions of $i < j$ where $i > j_0$. Next we define an approximation a_j to α_j. We let $a_j = |\{t < v_j : t \notin V_j\}|$ for all $j \ge 1$, where $V_j = \{v_1, v_2, \ldots, v_j\}$. Observe that $\alpha_j \le a_j$ for $j \le j_0$. Moreover,

$$\mathbf{Pr}(a_j = k) = (1-p_1)^k p_1 \quad \text{for } k \ge 0. \tag{12}$$

To see this, observe that the vertex v_j was chosen as the leftmost vertex available to the algorithm, and determining this vertex involves querying edges which have not yet been conditioned by the running of the algorithm. Observe that (12) holds even when conditioning on any previous history of the algorithm.

So $a_1 = 0$ and $a_2, a_3, \ldots$ is a sequence of independent copies of $Geo(p_1) - 1$ where $Geo(p_1)$ is the geometric random variable with probability of success p_1. We thus have:

$$\mathbf{E}(\mathrm{inv}(\sigma)) \le \mathbf{E}\left(\sum_{j=0}^{j_0} \alpha_j\right) \le \mathbf{E}\left(\sum_{j=0}^{j_0} a_j\right) \le \mathbf{E}\left(\sum_{j=0}^{n} a_j\right) \le n\frac{1-p_1}{p_1}. \tag{13}$$

Moreover, standard concentration arguments give then that

$$\mathrm{inv}(\sigma) \le \sum_{j=1}^{j_0} \alpha_j \le \sum_{j=1}^{j_0} a_j \le \sum_{j=1}^{n} a_j \le \frac{n}{p_1} \; q.s. \tag{14}$$

It follows from Lemma 3.3, (13), and (14) that

$$\mathbf{E}(\Lambda_1) \le n\left(\frac{4}{p_1} - 3\right) \text{ and } \Lambda_1 \le n\left(1 + \frac{4}{p_1}\right) \; q.s. \tag{15}$$

It remains to show that there is a Hamilton cycle of length not much greater then Λ_1. We show that we can cheaply insert all $\sigma_j, j > j_0$ into P_1 by replacing edges by paths of length two.

We let $I = \left\{i \le j_0 : |v_i - v_{i-1}|, |v_i - v_{i+1}| \le n^{2/3}\right\}$ and $m = |I|$. It follows from (15) that $m \ge n_0 - O(n^{1/3})$ q.s. For $J = \{\sigma_{j_0+1}, \ldots, \sigma_n\}$, our aim is to use the edges of G_2 to insert J into the path $v_1, \ldots, v_{j_0} = \sigma_1, \ldots, \sigma_{j_0}$, thus extending it to all of G_1.

Let $v_j \in J$ be a vertex to be inserted. Without loss of generality, assume $v_j \ge n/2$; the reverse case is essentially the same. We proceed by examining $v_j - 1, v_j - 2, v_j - 3, \ldots$ until we find a ℓ satisfying

(i) $\ell \in I$

(ii) $\{v_\ell, v_j\}, \{v_{\ell-1}, v_j\} \in E(G_2)$.

When examining a vertex v_ℓ for which ℓ satisfies (i), there is a p_1^2 chance that v_ℓ satisfies (ii). In particular, after examining $\log^2 n$ vertices $v = v_\ell$ which satisfy (i), we will q.s. find a v which satisfies both (i) and (ii). Thus there is q.s. a vertex $v_\ell \leq v_j$ satisfying (i) and (ii) with $v_j - v_\ell \leq \log^2 n + O(n^{1/3})$. Thus, we can replace the edge $(v_{\ell-1}, v_\ell)$ by a path $v_{\ell-1}, v_j, v_\ell$ to q.s. incorporate v_j into our path at a cost of at most $O\left(\log^2 n + n^{1/3} + n^{2/3}\right)$. Finally, (11) implies that we can insert all vertices of J in this manner one-by-one (even when we avoid re-using a candidate v_ℓ, this only increases the patching distance by at most a $\frac{\log^2 n}{p_1}$ additive factor). In the end, we get a Hamilton path $x_1, x_2, \ldots, x_n$ in $G_1 \cup G_2$ q.s. (with $v_1 \ldots, v_{j_0}$ a subsequence of $x_1, \ldots, x_n$) and the total added cost over Λ_1 is q.s. $O(n^{2/3} \log^2 n)$. There is only an exponentially small probability that we cannot find G_3-edges $\{x_1, x_{j+1}\}$, $\{x_j, x_n\}$ which now give us a Hamilton cycle; since the maximum value of of Z is just n^2, this gives $\mathbf{E}(Z) \leq A_p n$, as desired. □

We have:

Corollary 3.5. *Suppose that we replace the length of edge (i, j) in Lemma 3.4 by $\xi_i + \cdots + \xi_{j-1}$ where $\xi_1, \xi_2, \ldots, \xi_n$ are random variables with mean bounded above by μ and exponential tails. If $\xi_1, \ldots, \xi_n$ are independent of $G_{n,p}$ then $\mathbf{E}(Z) \leq A_p \mu n$.*

Proof. The bound on the expectation follows directly from Lemma 3.4 and the linearity of expectation. □

Let us observe now that we get an upper bound $\mathbf{E}(T(\mathcal{Y}_{t,p}^1)) \leq A_p t$ on the length of a tour in 1 dimension. We have

$$\mathbf{E}(T(\mathcal{Y}_{t,p}^1)) = \sum_{n=0}^{\infty} \mathbf{E}\left(T(\mathcal{Y}_{t,p}^1) \big| |\mathcal{Y}_{t,p}^1| = n\right) \mathbf{Pr}(|\mathcal{Y}_{t,p}^1| = n).$$

When conditioning on $|\mathcal{Y}_{t,p}^1| = n$, we let $P_1 < P_2 < \cdots < P_n \subset [0, t]$ be the points in $\mathcal{Y}_{t,p}^1$. We choose $k \in \{0, n-1\}$ uniformly randomly and let $\xi_i = ||P_{k+i+1} - P_{k+i}||$, where the indices of the P_j are evaluated modulo n. We now have $\mathbf{E}(\xi_i) \leq \frac{2t}{n}$ for all i, and Corollary 3.5 gives that

$$\mathbf{E}\left(T(\mathcal{Y}_{t,p}^1) \big| |\mathcal{Y}_{t,p}^1| = n\right) \leq A_p n \cdot \frac{2t}{n},$$

and thus

$$\mathbf{E}\left(T(\mathcal{Y}_{t,p}^1)\right) \leq 2A_p t. \tag{16}$$

3.2 The asymptotic tour length

Our proof of Theorem 3.1 uses recursion, by dividing the $[t]^d$ cube into smaller parts. However, since our divisions of the cube must not cross boundaries of the elemental regions $\mathcal{Y}_1^d$, we cannot restrict ourselves to subdivisions into perfect cubes (in general, the integer t may not have the divisors we like).

To this end, if $L = T_1 \times T_2 \times \cdots \times T_d$ where each T_i is either $[0,t]$ or $[0,t-1]$, we say L is a d-dimensional *near-cube* with sidelengths in $\{t-1,t\}$. For $0 \le d' \le d$, we define the canonical example $L_d^{d'} := [0,t]^{d'} \times [0,t-1]^{d-d'}$ for notational convenience, and let

$$\Phi_p^{d,d'}(t) = \mathbf{E}\left(T(\mathcal{Y}_{t,p}^d \cap L_d^{d'})\right).$$

so that

$$\Phi_p^d(t) := \Phi_p^{d,d}(t) = \Phi_p^{d,0}(t+1).$$

In the unlikely event that $\mathcal{Y}_{t,p}^d \cap L_d^{d'}$ is not Hamiltonian, we take $T(\mathcal{Y}_{t,p}^d \cap L_d^{d'}) = t^{d+1}\sqrt{d}$, for technical reasons.

Our first goal is an asymptotic formula for Φ:

Lemma 3.6. *There exists $\beta = \beta_{p,d} > 0$ such that*

$$\Phi_p^{d,d'}(t) \approx \beta t^d.$$

The proof of this is deferred until after the proof of Corollary 3.9 below.

The proof is by induction on $d \ge 2$. We prove the base case $d = 2$ along with the general case. We begin with a technical lemma.

Lemma 3.7. *For every fixed $p, d \ge d' \ge 0$, there is a constant $F_{p,d} > 0$ such that*

$$\Phi_p^{d,d'}(t) \le \Phi_p^{d,d'-1}(t) + F_{p,d}t^{d-1} \tag{17}$$

for all t sufficiently large. Here we interpret $\Phi_p^{d,d'}(t) = \Phi_p^{d,0}(t)$ for $d' < 0$. In particular, this implies that there is a constant $A_{p,d} > 0$ such that

$$\Phi_p^d(t+h) \le \Phi_p^d(t) + A_{p,d}ht^{d-1} \tag{18}$$

for sufficiently large t and $1 \le h \le t$.

Proof. We let S denote the subgraph of $\mathcal{Y}_{t,p}^d \cap L_d^{d'}$ induced by the difference $L_d^{d'} \setminus L_d^{d'-1}$.

By ignoring the d'th coordinate of S if $d' > 0$ and the dth coordinate otherwise, we obtain the $(d-1)$ dimensional set $\pi(S)$, for which induction on d (or equation (16) if $d = 2$) implies an expected tour $T(S)$ of length $\Phi_p^{d-1,d'-1}(t) \le \beta_p^{d-1}t^{d-1}$, and so changing notation, we can write

$$\Phi_p^{d-1,d'-1}(t) \le D_{p,d-1}t^{d-1}.$$

We have that

$$\mathbf{E}(T(S)) \le \mathbf{E}(T(\pi(S)) + d^{1/2}\,\mathbf{E}(|\pi(S)|) \le D_{p,d-1}t^{d-1} + d^{1/2}t^{d-1}.$$

The first inequality stems from the fact that the points in $L_d^{d'} \setminus L_d^{d'-1}$ have a d' coordinate in $[t-1,t]$.

Now if $\mathcal{Y}_{t,p}^d \cap L_d^{d'-1}$ and S are both Hamiltonian, then we have

$$T(\mathcal{Y}_{t,p}^d \cap L_d^{d'}) \le T(\mathcal{Y}_{t,p}^d \cap L_d^{d'-1}) + T(S) + O_d(t) \tag{19}$$

which gives us the Lemma, by linearity of expectation. We have (19) because we can patch together the minimum cost Hamilton cycle H in $\mathcal{Y}_{t,p}^d \cap L_d^{d'-1}$ and the minimum cost path P in S as follows: Let u_1, v_1 be the endpoints of P. If there is an edge u, v of H such that $(u_1, u), (v_1, v)$ is an edge in $\mathcal{Y}_{t,p}^d$ then we can create a cycle H_1 through $\mathcal{Y}_{t,p}^d \cap L_d^{d'-1} \cup P$ at an extra cost of at most $2d^{1/2}t$. The probability there is no such edge is at most $(1-p^2)^{t/2}$, which is negligible given the maximum value of $T(\mathcal{Y}_{t,p}^d \cap L_d^{d'})$.

On the other hand, because p is a constant, the probability that either of $\mathcal{Y}_{t,p}^d \cap L_d^{d'-1}$ or S is not Hamiltonian is exponentially small in t, (see for example [9]), which is again negligible given the maximum value of $T(\mathcal{Y}_{t,p}^d \cap L_d^{d'})$. This completes the proof of (17).

To obtain (18) we use (17) to write

$$\Phi_p^{d,d}(t+h) \leq \Phi_p^{d,0}(t+h) + dF_{p,d}(t+h)^{d-1} = \Phi_p^d(t+h-1) + dF_{p,d}(t+h)^{d-1} \leq \Phi_p^d(t) + dF_{p,d}\sum_{i=0}^{h}(t+i)^{d-1}.$$

□

Our argument is an adaptation of that in Beardwood, Halton and Hammersley [2] or Steele [21], with modifications to address difficulties introduced by the random set of available edges. First we introduce the concept of a decomposition into near-cubes. (Allowing near-cube decompositions is necessary for the end of the proof, beginning with Lemma 3.10). Simplifications relying on *Boundary Functionals* as in Yukich [22] do not appear to be available due to missing edges.

We say that a partition of $L_d^{d'}$ into m^d near-cubes S_α with sidelengths in $\{u, u+1\}$ indexed by $\alpha \in [m]^d$ is a *decomposition* if for each $1 \leq b \leq d$, there is an integer M_b such that, letting

$$f_b(a) = \begin{cases} au \text{ if } a < M_b \\ (a - M_b)(u+1) + M_b u \text{ if } a \geq M_b. \end{cases}.$$

we have that

$$S_\alpha = [f_1(\alpha_1 - 1), f_1(\alpha_1)] \times [f_2(\alpha_2 - 1), f_2(\alpha_2)] \times \cdots \times [f_d(\alpha_d - 1), f_d(\alpha_d)].$$

Observe that so long as $u \ll t$, $L_d^{d'}$ always has a decomposition into near-cubes with sidelengths in $\{u, u+1\}$. Indeed, if $t = ru - s$ for $0 \leq s < u$ then we can take $M_b = s$ for $b \leq d'$ and $M_b = s - 1$ for $b > d'$, unless $s = 0$, in which case $M_b = u - 1$.

First we note that tours in not-too-small near-cubes of a decomposition can be pasted together into a large tour at a reasonable cost:

Lemma 3.8. *Fix $\delta > 0$, and suppose $t = mu$ for $u = t^\gamma$ for $\delta < \gamma \leq 1$ $(m, u \in \mathbb{Z})$, and suppose S_α $(\alpha \in [m]^d)$ is a decomposition of $L_d^{d'}$. We let $\mathcal{Y}_{t,p}^{d,\alpha} := \mathcal{Y}_{t,p}^d \cap S_\alpha$. We have*

$$T(\mathcal{Y}_{t,p}^d \cap L_d^{d'}) \leq \sum_{\alpha \in [m]^d} T(\mathcal{Y}_{t,p}^{d,\alpha}) + 4m^d u\sqrt{d} \qquad \textit{with probability at least} \quad 1 - e^{-\Omega(u^d p^2)}.$$

Proof. Let $\mathcal{B}, \mathcal{C}$ denote the events

$$\mathcal{B} = \left\{\exists \alpha : \mathcal{Y}_{t,p}^{d,\alpha} \text{ is not Hamiltonian}\right\}$$

$$\mathcal{C} = \left\{ \exists \alpha : \left| |\mathcal{Y}_{t,p}^{d,\alpha}| - u^d \right| \geq \delta u^d \right\},$$

and let $\mathcal{E} = \mathcal{B} \cup \mathcal{C}$.

Now $\mathbf{Pr}(\mathcal{B}) \leq m^d e^{-\Omega(u^d p)}$ and, by Observation 3.2, $\mathbf{Pr}(\mathcal{C}) \leq m^d e^{-\Omega(u^d)}$ and so $\mathbf{Pr}(\mathcal{E}) \leq e^{-\Omega(u^d p)}$. Assume therefore that $\neg\mathcal{E}$ occurs. Each subcube S_α will contain a minimum length tour H_α. We now order the subcubes $\{S_\alpha\}$ as $T_1, \ldots, T_{m^d}$, such that for $S_\alpha = T_i$ and $S_{\alpha'} = T_{i+1}$, we always have that the Hamming distance between α and α' is 1. Our goal is to inductively assemble a tour through the subcubes $T_1, T_2, \ldots, T_j$ from the smaller tours H_α with a small number of additions and deletions of edges.

Assume inductively that for some $1 \leq j < m^d$ we have added and deleted edges and found a single cycle C_j through the points in $T_1, \ldots, T_j$ in such a way that (i) the added edges have total length at most $4\sqrt{d}ju$ and (ii) we delete one edge from $\tau(T_1)$, $\tau(T_j)$ and two edges from each $\tau(T_i), 2 \leq i \leq j-1$. To add the points of T_{j+1} to create C_{j+1} we delete one edge (u,v) of $\tau(T_j) \cap C_j$ and one edge (x,y) of $\tau(T_{j+1})$ such that both edges $\{u,x\}, \{v,y\}$ are in the edge set of $\mathcal{Y}_{t,p}^d$. Such a pair of edges will satisfy (i) and (ii) and the probability we cannot find such a pair is at most $(1-p^2)^{(u^d/2-1)u^d/2}$. Thus with probability at least $1 - e^{\Omega(u^d p^2)}$ we build the cycle C_{m^d} with a total length of added edges $\leq 4\sqrt{d}m^d u$. □

Linearity of expectation (and the upper bound $t^{d+1}\sqrt{d}$ on $T(\mathcal{Y}_{t,p}^d)$ when there is no tour) now gives a short-range recursive bound on $\Phi_p^d(t)$ when t factors reasonably well:

Corollary 3.9. *For all large u and $1 \leq m \leq u^{10}$ $(m, u \in \mathcal{N})$,*

$$\Phi_p^d(mu) \leq m^d(\Phi_p^d(u) + B_{p,d}u)$$

for some constant B_d. □

Proof of Lemma 3.6.
Note that here we are using a decomposition of $[mu]^d$ into m^d subcubes with sidelength u; near-cubes are not required.

To get an asymptotic expression for $\Phi_p^d(t)$ we now let

$$\beta = \beta_{p,d} = \liminf_t \frac{\Phi_p^d(t)}{t^d}.$$

Choose u_0 large and such that

$$\frac{\Phi_p^d(u_0)}{u_0^d} \leq \beta + \varepsilon$$

and then define the sequence $u_k, k \geq -1$ by $u_{-1} = u_0$ and $u_{k+1} = u_k^{10}$ for $k \geq 0$. Assume inductively that for some $i \geq 0$ that for $A_{p,d}$ as in Lemma 3.7 and $B_{p,d}$ as in Corollary 3.9,

$$\frac{\Phi_p^d(u_i)}{u_i^d} \leq \beta + \varepsilon + \sum_{j=-1}^{i-2} \left(\frac{A_{p,d}}{u_j} + \frac{B_{p,d}}{u_j^{d-1}} \right). \tag{20}$$

This is true for $i = 0$, and then for $i \geq 0$ and $0 \leq u \leq u_i$ and $d \leq m \in [u_{i-1}, u_{i+1}]$ we have

$$\frac{\Phi_p^d(mu_i + u)}{(mu_i + u)^d} \leq \frac{\Phi_p^d(mu_i) + A_{p,d}u(mu_i)^{d-1}}{(mu_i)^d}, \qquad \text{from Lemma 3.7,} \tag{21}$$

$$\leq \frac{m^d(\Phi_p^d(u_i) + B_{p,d}u_i) + A_{p,d}u(mu_i)^{d-1}}{(mu_i)^d}, \qquad \text{from Corollary 3.9,}$$

$$\leq \beta + \varepsilon + \sum_{j=-1}^{i-2}\left(\frac{A_{p,d}}{u_j} + \frac{B_{p,d}}{u_j^{d-1}}\right) + \frac{B_{p,d}}{u_i^{d-1}} + \frac{A_{p,d}}{m}, \qquad \text{by induction,}$$

$$\leq \beta + \varepsilon + \sum_{j=-1}^{i-1}\left(\frac{A_{p,d}}{u_j} + \frac{B_{p,d}}{u_j^{d-1}}\right). \tag{22}$$

Putting $m = u_{i+1}/u_i$ and $u = 0$ into (21) and (22) completes the induction. We deduce from (20), (21) and (22) that for $i \geq 0$ we have

$$\frac{\Phi_p^d(t)}{t^d} \leq \beta + \varepsilon + \sum_{j=-1}^{\infty}\left(\frac{A_{p,d}}{u_j} + \frac{B_{p,d}}{u_j^{d-1}}\right) \leq \beta + 2\varepsilon \qquad \text{for } t \in J_i = [u_{i-1}u_i, u_i(u_{i+1}+1)]$$

Now $\bigcup_{i=0}^{\infty} J_i = [u_0^2, \infty]$ and since ε is arbitrary, we deduce that

$$\beta = \lim_{t\to\infty} \frac{\Phi_p^d(t)}{t^d}, \tag{23}$$

We can conclude that

$$\Phi_p^d(t) \approx \beta t^d,$$

which, together with Lemma 3.7, completes the proof of Lemma 3.6, once we show that $\beta > 0$ in (23). To this end, we let ρ denote $\mathbf{Pr}(|\mathcal{Y}_1^d| \geq 1)$, so that $\mathbf{E}(|\mathcal{Y}_t^d|) \geq \rho t^d$. We say $x \in \{0, \ldots, t-1\}^d$ is *occupied* if there is a point in the copy $\mathcal{Y}_1^d + x$. Observing that a unit cube $[0,1]^d + x$ $(x \in \{0, \ldots, t-1\}^d)$ is at distance at least 1 from all but $3^d - 1$ other cubes $[0,1]^d + y$, we certainly have that the minimum tour length through $\mathcal{Y}_t^d$ is at least $\frac{\mathcal{O}}{3^d-1}$, where where $\mathcal{O}$ is the number of occupied x. Linearity of expectation now gives that $\beta > \rho/(3^d - 1)$, completing the proof of Lemma 3.6. □

Before continuing, we prove the following much cruder version of Part (ii) of Theorem 3.1:

Lemma 3.10. *For any fixed $\varepsilon > 0$, $T(\mathcal{Y}_{t,p}^d) \leq t^{d+\varepsilon}$ q.s.*

Proof. We let $m = \lfloor t^{1-\varepsilon/2} \rfloor$, $u = \lfloor t/m \rfloor$, and let $\{\mathcal{Y}_{\tau,p}^{d,\alpha}\}$ be a decomposition of $\mathcal{Y}_{t,p}^d$ into m^d near-cubes with sidelengths in $\{u, u+1\}$. We have that q.s. each $\mathcal{Y}_{\tau,p}^{d,\alpha}$ has (i) $\approx u^d$ points, and (ii) a Hamilton cycle H_α. We can therefore q.s. bound all $T(\mathcal{Y}_{\tau,p}^{d,\alpha})$ by $du \cdot u^d$, and Lemma 3.8 gives that q.s. $T(\mathcal{Y}_{t,p}^d) \leq 4dut^d + 4m^d u\sqrt{d}$. □

Proof of Theorem 3.1.
We consider a decomposition $\{S_\alpha\}$ $(\alpha \in [m]^d)$ of $\mathcal{Y}_t^d$ into m^d near-cubes of side-lengths in $\{u, u+1\}$, for $\gamma = 1 - \frac{\varepsilon}{2}$, $m = \lfloor t^\gamma \rfloor$, and $u = \lfloor t/m \rfloor$.

Lemma 3.6 gives that

$$\mathbf{E}\, T(\mathcal{Y}_{t,p}^{d,\alpha}) \approx \beta u^d \approx \beta t^{(1-\gamma)d}.$$

Let

$$\mathcal{S}_\gamma(\mathcal{Y}_{t,p}^d) = \sum_{\alpha \in [m]^d} \min\left\{T(\mathcal{Y}_{t,p}^{d,\alpha}), 2dt^{(1-\gamma)(d+\varepsilon)}\right\}.$$

Note that $\mathcal{S}_\gamma(\mathcal{Y}_{t,p}^d)$ is the sum of $t^{\gamma d}$ identically distributed bounded random variables.

Now, since q.s. $T(\mathcal{Y}_{t,p}^{d,\alpha}) \leq 2dt^{(1-\gamma)(d+\varepsilon)}$ for all α by Lemma 3.10, we have that q.s. $\mathcal{S}_\gamma(\mathcal{Y}_{t,p}^d) = \sum_\alpha T(\mathcal{Y}_{t,p}^{d,\alpha})$. Applying Theorem 1 of Hoeffding [11] for the sum of independent bounded random variables, we see that for any $\xi > 0$, we have

$$\mathbf{Pr}(|\mathcal{S}_\gamma(\mathcal{Y}_{t,p}^d) - m^d\,\mathbf{E}(T(\mathcal{Y}_{u,p}^d))| \geq \xi) \leq 2\exp\left(-\frac{2\xi^2}{4m^d d^2 t^{2(1-\gamma)(d+\varepsilon)}}\right).$$

Putting $\xi = t^{d\varepsilon}$ for small ε, we see that

$$\mathcal{S}_\gamma(\mathcal{Y}_{t,p}^d) = \beta t^d + o(t^d) \qquad q.s. \tag{24}$$

Note next that Lemma 3.8 implies that

$$T(\mathcal{Y}_{t,p}^d) \leq \mathcal{S}_\gamma(\mathcal{Y}_{t,p}^d) + \delta_2 \text{ where } \delta_2 = o(t^d) \qquad q.s. \tag{25}$$

It follows from (24) and (25) and the fact that $\mathbf{Pr}(|\mathcal{Y}_t^d| = t^d) = \Omega(t^{-d/2})$ that

$$T(\mathcal{Y}_{t,p}^d) \leq \beta t^d + o(t^d) \qquad q.s. \tag{26}$$

which proves part (ii) of Theorem 3.1.

Of course, we have from Lemma 3.6 that

$$\mathbf{E}(T(\mathcal{Y}_{t,p}^d)) = \beta t^d + \delta_1 \text{ where } \delta_1 = o(t^d), \tag{27}$$

and we show next that that this together with (25) implies part (i) of Theorem 3.1, that:

$$T = T(\mathcal{Y}_{t,p}^d) = \beta t^d + o(t^d) \qquad a.a.s. \tag{28}$$

We choose $0 \leq \delta_3 = o(t^d))$ such that $0 \leq \delta_2, |\delta_1| = o(\delta_3)$. Let $I = [\beta t^d - \delta_3, \beta t^d + \delta_2]$. Then we have

$$\begin{aligned}\beta t^d + \delta_1 = \mathbf{E}(T(\mathcal{Y}_{t,p}^d) \mid T(\mathcal{Y}_{t,p}^d) \geq \beta t^d + \delta_2)\,\mathbf{Pr}(T(\mathcal{Y}_{t,p}^d) \geq \beta t^d + \delta_2)\\ + \mathbf{E}(T(\mathcal{Y}_{t,p}^d) \mid T(\mathcal{Y}_{t,p}^d) \in I)\,\mathbf{Pr}(T(\mathcal{Y}_{t,p}^d) \in I) +\\ \mathbf{E}(T(\mathcal{Y}_{t,p}^d) \mid T(\mathcal{Y}_{t,p}^{d,\alpha}) \leq \beta t^d - \delta_3)\,\mathbf{Pr}(T(\mathcal{Y}_{t,p}^d) \leq \beta t^d - \delta_3).\end{aligned}$$

Now $\varepsilon_1 = \mathbf{E}(T(\mathcal{Y}_{t,p}^d) \mid T(\mathcal{Y}_{t,p}^d) \geq \beta t^d + \delta_2)\,\mathbf{Pr}(T(\mathcal{Y}_{t,p}^d) \geq \beta t^d + \delta_2) = O(t^{-\omega(1)})$ since $|\mathcal{Y}_{t,p}^d| \leq 2d^{1/2}t^d$ and $\mathbf{Pr}(T(\mathcal{Y}_{t,p}^d) \geq \beta t^d + \delta_2) = O(t^{-\omega(1)})$, from (26).

So, if $\lambda = \mathbf{Pr}(T(\mathcal{Y}_{t,p}^d) \in I)$ then we have

$$\beta t^d + \delta_1 \leq \varepsilon_1 + (\beta t^d + \delta_2)\lambda + (\beta t^d - \delta_3)(1-\lambda)$$

or

$$\lambda \geq \frac{\delta_1 - \varepsilon_1 + \delta_3}{\delta_2 + \delta_3} = 1 - o(1),$$

and this proves (28) completing the proof of Theorem 3.1. □

Proof of Theorem 1.3.
We now let $\mathcal{W}^d_{t,p}$ be the graph on the set of points in $[0,t]^d$ which is the result of a Poisson process of intensity 1. Our first task is to bound the variance $\mathcal{V}(t)$ of $T(\mathcal{W}^d_{t,p})$. Here we follow Steele's argument [21] with only small modifications. We approximate $T(\mathcal{W}^d_{2t,p})$ as the sum over 2^d half-size cubes of $T(\mathcal{W}^d_{t,p})$ and use this to show that $\sum_{k=1}^{\infty} \frac{\mathcal{V}(2^k t)}{(2^k t)^{2d}} \leq \infty$. This deals with n of the form $2^k t$ for some value of t and we then have to fill in the gaps.

Let $\mathcal{E}_t$ denote the event that

$$T(\mathcal{W}^d_{2t,p}) \leq \sum_{\alpha \in [2]^d} T(\mathcal{W}^{d,\alpha}_{t,p}) + 2^{d+2} t\sqrt{d}. \tag{29}$$

Observe that Lemma 3.8 with $m = 2, u = t/2$ implies that

$$\mathbf{Pr}(\neg \mathcal{E}_t) \leq e^{-\Omega(t^d p)}. \tag{30}$$

We define the random variable $\lambda(t) = T(\mathcal{W}^d_{t,p}) + 10t\sqrt{d}$, and let λ_i denote independent copies of $\lambda(t)$. Conditioning on $\mathcal{E}_t$, we have from (29) that

$$\lambda(2t) \leq \sum_{i=1}^{2^d} \lambda_i(t) - 6t2^d\sqrt{d} \leq \sum_{i=1}^{2^d} \lambda_i(t).$$

In particular, (30) implies that letting $\Upsilon(t) = \mathbf{E}(\lambda(t)) = \Omega(t^d)$ (see (27)) and $\Psi(t) = \mathbf{E}(\lambda(t)^2)$, we have for sufficiently large t that

$$\begin{aligned}
\Psi(2t) &\leq \mathbf{E}\left(\left(\sum_{\alpha \in [2]^d} T(\mathcal{W}^{d,\alpha}_{t,p}) + 2^{d+2}t\sqrt{d} + 21t\sqrt{d}\right)^2\right) \\
&= \sum_{i=1}^{2^d} \mathbf{E}((\lambda_i(t) - 10t\sqrt{d})^2) + \sum_{i \neq j}^{2^d} \mathbf{E}(\lambda_i(t) - 10t\sqrt{d})\,\mathbf{E}(\lambda_j(t) - 10t\sqrt{d}) + \\
&\qquad + (2^{d+2} + 21)t\sqrt{d} \sum_{i=1}^{2^d} \mathbf{E}(\lambda_i(t) - 10t\sqrt{d}) + ((2^{d+2} + 21)t\sqrt{d})^2 \\
&= 2^d\,\mathbf{E}((\lambda(t) - 10t\sqrt{d})^2) + 2^d(2^d - 1)\,\mathbf{E}(\lambda(t) - 10t\sqrt{d})^2 + \\
&\qquad + 2^d(2^{d+2} + 21)t\sqrt{d}\,\mathbf{E}(\lambda(t) - 10t\sqrt{d}) + ((2^{d+2} + 21)t\sqrt{d})^2 \\
&= 2^d\Psi(t) + 2^d(2^d - 1)\Upsilon(t)^2 - \Omega(t\,\mathbf{E}(\lambda(t)) + O(t^2)) \\
&\leq 2^d\Psi(t) + 2^d(2^d - 1)\Upsilon(t)^2.
\end{aligned}$$

For

$$\mathcal{V}(t) := \mathbf{Var}(T(\mathcal{W}^d_{t,p})) = \Psi(t) - \Upsilon(t)^2,$$

we have

$$\frac{\mathcal{V}(2t)}{(2t)^{2d}} - \frac{1}{2^d}\frac{\mathcal{V}(t)}{t^{2d}} \leq \frac{\Upsilon(t)^2}{t^{2d}} - \frac{\Upsilon(2t)^2}{(2t)^{2d}}.$$

Now with $t \geq 1$ arbitrary, summing over $2^k t$ for $k = 0, \ldots, M-1$ gives

$$\sum_{k=1}^{M} \frac{\mathcal{V}(2^k t)}{(2^k t)^{2d}} - \frac{1}{2^d} \sum_{k=0}^{M-1} \frac{\mathcal{V}(2^k t)}{(2^k t)^{2d}} \leq \frac{\Upsilon(t)^2}{t^{2d}} - \frac{\Upsilon(2^M t)^2}{(2^M t)^{2d}} \leq \frac{\Upsilon(t)^2}{t^{2d}}$$

and so, solving for the first sum, we find

$$\sum_{k=1}^{M} \frac{\mathcal{V}(2^k t)}{(2^k t)^{2d}} \leq \left(1 - \frac{1}{2^d}\right)^{-1} \left(\frac{\mathcal{V}(t)}{t^{2d}} + \frac{\Upsilon(t)^2}{t^{2d}}\right) < \infty. \tag{31}$$

Still following Steele, we let $N(t)$ be the Poisson counting process on $[0,\infty)$. We fix a random embedding $\mathcal{U}$ of $\mathcal{N}$ in $[0,1]^d$ as $u_1, u_2, \ldots$ and a random graph $\mathcal{U}_p$ where each edge is included with independent probability p. We let $\mathcal{U}_{n,p}$ denote the restriction of this graph to the first n natural numbers. In particular, note that $\mathcal{U}_{N(t^d),p}$ is equivalent to $\mathcal{W}_{t,p}$, scaled from $[0,t]^d$ to $[0,1]^d$. Thus, applying Chebychev's inequality to (31) gives, in conjunction with Lemma 3.6, that

$$\sum_{k=0}^{\infty} \mathbf{Pr}\left(\left|\frac{t2^k T(\mathcal{U}_{N((t2^k)^d),p})}{(t2^k)^d} - \beta_{p,d}\right| > \varepsilon\right) < \infty$$

and so for $t > 0$ that

$$\lim_{k\to\infty} \frac{T(\mathcal{U}_{N((t2^k)^d),p})}{(t2^k)^{d-1}} = \beta_{p,d} \qquad a.s. \tag{32}$$

Now choosing some large integer ℓ, we have that (32) holds simultaneously for all the (finitely many) integers $t \in S_P = [2^\ell, 2^{\ell+1})$; and for $2^\ell \leq r \in \mathbb{R}$, we have that

$$r \in [2^k t, 2^k(t+1)) \text{ for } t \in S_\ell \text{ and some } k. \tag{33}$$

(We simply choose k such that $2^\ell \leq 2^{-k} r < 2^{\ell+1}$.)

Unlike the classical case $p = 1$, in our setting, we do not have monotonicity of $T(\mathcal{U}_{n,p})$. Nevertheless, we show a kind of continuity of the tour length through $T(\mathcal{U}_{n,p})$:

Lemma 3.11. *For all $\varepsilon > 0$, $\exists \delta > 0$ such that for all $0 \leq k < \delta n$, we have*

$$T(\mathcal{U}_{n+k,p}) < T(\mathcal{U}_{n,p}) + \varepsilon n^{\frac{d-1}{d}}, \qquad q.s. \tag{34}$$

Proof. We consider cases according to the size of k.

Case 1: $k \leq n^{\frac{1}{3}}$.
Note that we have $T(\mathcal{U}_{n+1,p}) < T(\mathcal{U}_{n,p}) + \sqrt{d}$ q.s., since we can q.s. find an edge in the minimum tour though $\mathcal{U}_{n,p}$ whose endpoints are both adjacent to $(n+1)$. $n^{\frac{1}{3}}$ applications of this inequality now give (34).

Case 2: $k > n^{\frac{1}{3}}$.
In this case the restriction $\mathcal{R}$ of $\mathcal{U}_{n+k,p}$ to $\{n+1,\ldots,k\}$ is q.s. (with respect to n) Hamiltonian [4]. In particular, by Theorem 3.1, we can q.s. find a tour T though $\mathcal{R}$ of length $\leq 2\beta_p^d k^{\frac{d-1}{d}}$. Finally, there are q.s., edges $\{x,y\}$ and $\{w,z\}$ on the minimum tours through $\mathcal{U}_{n,p}$ and $\mathcal{R}$, respectively, such that $x \sim w$ and $y \sim z$ in $\mathcal{U}_{n+k,p}$, giving a tour of length

$$T(\mathcal{U}_{n+k,p}) \leq T(\mathcal{U}_{n,p}) + 2\beta_{p,d} k^{\frac{d-1}{d}} + 4\sqrt{d}. \qquad \square$$

This gives the lemma with $\delta = (\varepsilon/3\beta)^{d/(d-1)}$.

Let $\varepsilon_\ell, \ell = 1, 2, \ldots$ be a sequence of positive reals, tending to aero. We apply Lemma 3.11 with $\varepsilon = \varepsilon_\ell$ and $\delta = (1+\frac{1}{t})^d - 1 = O(\frac{d}{t})$, assuming t is large. Then we have $(2^k t)^d \leq r^d \leq (2^k t)^d(1+\delta) = (2^k(t+1))^d$ by (33), and using the fact that

$$(1-2\delta)N(r^d) < N((1-\delta)r^d) < N((1+\delta)r^d) < (1+2\delta)N(r^d) \text{ } q.s. \text{ (with respect to } r),$$

gives that for large enough ℓ, we have q.s.

$$T(\mathcal{U}_{N(((t+1)2^k)^d),p}) - \varepsilon_\ell r^{d-1} < T(\mathcal{U}_{N(r^d),p}) < T(\mathcal{U}_{N((t2^k)^d),p}) + \varepsilon_\ell r^{d-1},$$

and so dividing by r^{d-1} and using (32) and taking limits we find that a.s.

$$\beta_{p,d} - 2\varepsilon_\ell \leq \liminf_{r\to\infty} \frac{T(\mathcal{U}_{N(r^d)})}{r^{d-1}} \leq \limsup_{r\to\infty} \frac{T(\mathcal{U}_{N(r^d)})}{r^{d-1}} \leq \beta_{p,d} + 2\varepsilon_\ell.$$

Since ℓ may be arbitrarily large, we find that

$$\lim_{r\to\infty} \frac{T(\mathcal{U}_{N(r^d)})}{r^{d-1}} = \beta_{p,d}.$$

Now the elementary renewal theorem guarantees that

$$N^{-1}(n) \approx n, \qquad a.s.$$

So we have a.s.

$$\lim_{r\to\infty} \frac{T(\mathcal{U}_{n,p})}{n^{\frac{d-1}{d}}} = \lim_{r\to\infty} \frac{T(\mathcal{U}_{N(N^{-1}(n)),p})}{(N^{-1}(n))^{\frac{d-1}{d}}} \frac{(N^{-1}(n))^{\frac{d-1}{d}}}{n^{\frac{d-1}{d}}} = \beta_{p,d} \cdot 1 = \beta_{p,d}.$$

□

3.3 The case $p(n) \to 0$

We will in fact show that (1) holds q.s. for $np \geq \omega \log n$, for some $\omega \to \infty$. That we also get the statement of Theorem 1.2 can be seen by following the proof carefully, but this also follows as a consequence directly from the appendix in Johannsson, Kahn and Vu [12].

We first show that q.s.

$$T(\mathcal{X}_{n,p}) = \Omega(n^{(d-1)/d}/p^{1/d}). \tag{35}$$

Let Y_1 denote the number of vertices whose closest $G_{n,p}$-neighbor is at distance at least $\frac{1}{(np)^{1/d}}$. Observe first that if $r = 1/(np)^{1/d}$ then with probability at least $\left(1 - \nu_d r^d p\right)^{n-1} \approx e^{-\nu_d}$, there are no points within distance $1/(np)^{1/d}$ of any fixed $v \in \mathcal{X}_{n,p}$. Thus $\mathbf{E}(Y_1) \geq ne^{-\nu_d}/2$ and one can use the Azuma-Hoeffding inequality to show that Y_1 is concentrated around its mean. Thus q.s. $T(\mathcal{X}_{n,p}) \geq n^{(d-1)/d}e^{-\nu_d}/4p^{1/d}$, proving (35).

We will for convenience prove the following theorem. After which Theorem 1.2 follows in a couple of lines.

Theorem 3.12. *Let $\mathcal{Y}_1^d \subset [0,t]^d$ denote a set of points chosen via a Poisson process of intensity one in $[0,t]^d$ where $t = n^{1/d}$. Then there exists a constant γ_p^d such that*

$$T(\mathcal{Y}_{t,p}^d) \le \gamma_p^d \frac{t^d}{p^{1/d}} \qquad q.s.$$

Proof. We let $p_0 = p_1 = p/3$ and $p_i = p_1/2^{i-1}, i = 2, \dots, k = \log_2 t$ and define p_{k+1} so that $1 - p = \prod_{j=1}^{k+1}(1 - p_j)$. We let $G_i = \mathcal{Y}_{t,p_i}^d$, $i = 0, 1, \dots, k+1$, where each G_i has the same vertex set $\mathcal{Y}_1^d$, but in which the edges are independently chosen. Observe that with this choice, we have that $\mathcal{Y}_{t,p}^d$ decomposes as $\mathcal{Y}_{t,p}^d = \bigcup_{i=0}^{k+1} G_i$.

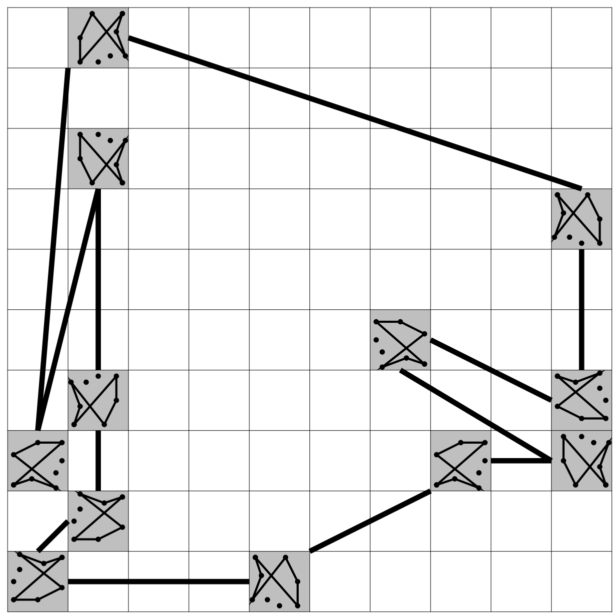

Figure 3: (Applying Theorem 3.1 to prove Theorem 3.12.) The shaded S_α are those which are heavy (with not too few vertices) and typical (with a large cycle). Theorem 3.1 implies that we can find the (short) cycle of bold edges through the S_α—these edges indicate the presence of *patchable pairs* which can be used to construct a large cycle. We then repeat this process on remaining vertices with a new set of random edges, with coarser and coarser divisions of the square, to cover most vertices with large cycles, which must then be merged.

Proof Strategy: We partition $[0,t]^d$ into small subcubes. Most of these will induce subgraphs of G_1 that contain relatively large cycles. We then use Theorem 3.1 to find a tour through these subcubes that can be used to patch together the subcube cycles into one large cycle H_1. We then use this idea and coarser and coarser partitions into subcubes to create cycles $H_2, H_3, \dots, H_\ell$ where we use G_j edges to create H_j. We then use the extensions and rotations of Pósa [18] and the edges of G_0 to merge the H_i and the relatively few vertices not in any of the H_i into a hamilton cycle of the required cost.

We begin by constructing a large cycle, using only the edges of G_1. We choose ε small and then choose an absolute constant K sufficiently large for subsequent claims. In preparation for an inductive argument we let $t_1 = t$, $T_1 = t_1^d$, $m_1 = \lfloor (T_1 p_1/K)^{1/d} \rfloor$ and consider the partition $\Delta_1 = \{S_\alpha\}$ $(\alpha \in [m_1]^d)$ of $[0,t]^d$ into m_1^d subcubes of side length $u = t/m_1$. (Note that t will not change throughout the induction).

Now each S_α contains $\approx K/p_1$ vertices, in expectation and so it has at least $(1-\varepsilon)K/p_1$ vertices with probability $1 - e^{-\Omega(K/p_1)} = 1 - o(1)$. Let α be *heavy* if S_α has at least this many vertices, and *light* otherwise. Let Γ_α be the subgraph of G_1 induced by S_α. If α is heavy then for any $\varepsilon > 0$ we can if K is sufficiently large find with probability at least $1 - e^{-\Omega(K/p_1)} = 1 - o(1)$, a cycle C_α in Γ_α containing at least $(1-\varepsilon)^2 K/p_1$ vertices. This is because when α is heavy, Γ_α has expected average degree at least $(1-\varepsilon)K$ (see [8] Chapter 6.3 for an explanation). We say that a heavy α is *typical* if it Γ_α contains a cycle with $(1-\varepsilon)|S_\alpha \cap \mathcal{X}|$ edges; otherwise it is *atypical*. We bound the length of aa C_α by $c_\alpha = |S_\alpha \cap \mathcal{X}| \times t_1 d^{1/2}/m_1$, where the second factor is the Euclidean diameter of S_α.

We now let N denote the set of vertices in $\bigcup C_\alpha$, where the union is taken over all typical heavy α. Our aim is to use Theorem 3.1(ii) to prove that we can q.s. merge the vertices N into a single cycle C_1, without too much extra cost, and using only the edges of G_1. Letting $q_\alpha = Pr(S_\alpha \text{ is typical}) \geq 1 - \varepsilon$, we make each typical heavy α *available* for this round with independent probability $\frac{1-\varepsilon}{1-q_\alpha}$, so that the probability that any given α is available is exactly $1 - \varepsilon$. (This is of course *rejection sampling.*) Now we can let $Y = \mathcal{Y}_1^d$ in Theorem 3.1 be a process which places a single point at the center of $[0,1]^d$ with probability $1-\varepsilon$, or produces an empty set with probability ε. Let now Y_α ($\alpha \in m_1^d$) be the independent copies of Y which give $\mathcal{Y}_{m_1}^d$. Given two cycles C_1, C_2 in a graph G we say that edges $u_i = (x_i, y_i) \in C_i, i = 1, 2$ are a *patchable pair* if $f_x = (x_1, x_2)$ and $f_y = (y_1, y_2)$ are also edges of G. Given $x \in Y_\alpha, y \in Y_\beta$, we let $x \sim y$ whenever there exist *two disjoint* patchable pairs $\sigma_{\alpha,\beta}$ between C_α, C_β. Observe that an edge between two vertices of $\mathcal{Y}_1^d$ is then present with probability

$$q_{\alpha,\beta} \geq \mathbf{Pr}(Bin((1-\varepsilon)^4 K^2/4p_1^2, p_1^2) \geq 2) \geq 1 - \varepsilon.$$

In particular, this graph contains a copy of $\mathcal{Y}_{m_1,(1-\varepsilon)}^d$, for which Theorem 3.1(ii) gives that q.s. we have a tour of length $\leq B_1 m_1^d$ for some constant B_1; in particular, there is a path $P = (\alpha_1, \alpha_2, \ldots, \alpha_M)$ through the typical heavy α with at most this length. Using P, we now merge its cycles $C_{\alpha_i}, i = 1, 2, \ldots, M$ into a single cycle.

Suppose now that we have merged $C_{\alpha_1}, C_{\alpha_2}, \ldots, C_{\alpha_j}$ into a single cycle C_j and have used one choice from $\sigma_{\alpha_{j-1},\alpha_j}$ to patch C_{α_j} into C_{j-1}. We initially had two choices for patching $C_{\alpha_{j+1}}$ into C_{α_j}, one may be lost, but one at least will be available. Thus we can q.s. use G_1 to create a cycle H_1 from $C_{\alpha_1}, C_{\alpha_2}$, by adding only patchable pairs of edges, giving a total length of at most

$$2T_1 \times \frac{t_1 d^{1/2}}{m_1} + m_1^d \left(B_1 + \frac{t_1 d^{1/2}}{m_1} \right) \leq \frac{L T_1 d^{1/2}}{p_1^{1/d}}, \tag{36}$$

where $L = 4K^{1/d}$.

The first term in (36) is a bound on the total length of the cycles C_α where α is available, assuming that $\sum_\alpha c_\alpha \leq |\mathcal{Y}_{t,p}^d| \leq 2t^d$. The second smaller term is the q.s. cost of patching these cycles to create H_1. This consists of the length of P plus a term bounding the cost of joining the center of a subcube C_α to the ends of an edge of C_α. This latter value, bounds the cost of adding C_α to H_1.

Having constructed H_1, we will consider coarser and coarser subdivisions $\mathcal{D}_i$ of $[0,t]^d$ into m_i^d subcubes, and argue inductively that we can q.s. construct, for each $1 \leq i \leq \ell$ for suitable ℓ, vertex disjoint cycles $H_1, H_2, \ldots, H_\ell$ satisfying:

P1 $T_i \leq 3\varepsilon T_{i-1}$ for $i \geq 2$, where $T_j = t^d - \sum_{i=1}^{j-1} |H_i|$,

P2 the set of points in the αth subcube in the decomposition $\mathcal{D}_i$ occupied by vertices which fail to participate in H_i is given by a process which occurs independently in each subcube in $\mathcal{D}_i$, and

P3 the total length of each H_i is at most $\frac{LT_i d^{1/2}}{p_i^{1/d}}$.

Note that H_1, above, satisfies these conditions for $\ell = 1$.

Assume inductively that we have constructed such a sequence $H_1, H_2, \dots, H_{j-1}$ ($j \geq 2$). We will now use the G_j edges to construct another cycle H_j. Suppose now that the set $\mathcal{T}_j$ of points that are not in $\bigcup_{i=1}^{j-1} H_i$ satisfies $T_j = |\mathcal{T}_j| \geq t^{d-1}/\log t$. We let $m_j = (T_j p_j/K)^{1/d}$ and $t_j = T_j^{1/d}$. The expected number of points in a subcube will be K/p_j but we have not exercised any control over its distribution. For $i \geq 2$, we let $\alpha \in [m_i]^d$ be heavy if S_α contains at least $\varepsilon K/p_j$ points. Now we want K to be large enough so that εK is large and that a heavy subcube has a cycle of size $(1-\varepsilon)|\mathcal{T}_j \cap S_\alpha|$ with probability at least $1-\varepsilon$, in which case, again, it is *typical.* We define Γ_j as the set of typical heavy pairs $\{\alpha, \beta\}$ for which there are at least two disjoint patchable pairs between the corresponding large cycles. Applying the argument above with T_j, t_j, m_j, Γ_j replacing T_1, t_1, m_1, Γ_1 (note that P2, above, ensures that Theorem 3.1 applies) we can q.s. find a cycle H_j with at least $(1-3\varepsilon)T_j$ vertices and length at most $\frac{LT_j d^{1/2}}{p_j^{1/d}}$, giving induction hypothesis part P3. Part P1 is satisfied since the light subcubes only contribute ε fraction of points to $\mathcal{T}_j$, and we q.s. take a $(1-\varepsilon)$ fraction of the heavy subcubes. Finally, Part P2 is satisfied since participation in H_j is determined exclusively by the set of adjacency relations in $G_j \cap \mathcal{T}_j$, which is independent of the positions of the vertices.

Thus we are guaranteed a sequence $H_1, H_2, \dots, H_\ell$ as above, such that $T_{\ell+1} < t^{d-1}/\log t$. The total length of $H_1, H_2, \dots, H_\ell$ is at most

$$\sum_{i=1}^{\ell} \frac{LT_i d^{1/2}}{p_i^{1/d}} \leq \frac{L3^{1/d} t^d}{p^{1/d}} \sum_{i=1}^{\infty} 3^i \cdot 2^{i/d} \varepsilon^{i-1} = O\left(\frac{t^d}{p^{1/d}}\right).$$

We can now use G_0 to finish the proof. It will be convenient to write $G_0 = \bigcup_{i=0}^{2} A_i$ where $A_i, i = 1, 2, 3$ are independent copies of $\mathcal{Y}_{t,q}^d$ where $1 - p_0 = (1-q)^3$. Also, let $R = \{x_1, x_2, \dots, x_r\} = \mathcal{Y}_{t,p}^d \setminus \bigcup_{i=1}^{\ell} H_i$.

We first create a Hamilton path containing all vertices, only using the edges of $A_1 \cup A_2$ and the extension-rotation algorithm introduced by Pósa [18]. We begin by deleting an arbitrary edge from H_1 to create a path P_1. Suppose inductively that we have found a path P_j through $Y_j = H_1 \cup \cdots H_{\rho_j} \cup X_j$, where $X_j \subseteq R$, at an added cost of $O(jt)$. We let V_j denote the vertices of P_j and promise that $V_{\ell+r} = \mathcal{Y}_{t,p}^d$. We also note that $|V_j| \geq |V_1| = \Omega(t^d)$ for $j \geq 1$.

At each stage of our process to create P_{j+1} we will construct a collection $\mathcal{Q} = \{Q_1, Q_2, \dots, Q_r\}$ of paths through V_j. Let $Z_{\mathcal{Q}}$ denote the set of endpoints of the paths in $\mathcal{Q}$. *Round* j of the process starts with P_j and is finished when we have constructed P_{j+1}.

If at any point in round j we find a path Q in $\mathcal{Q}$ with an endpoint x that is an A_2-neighbor of a vertex in $y \notin V_j$ then we will make a *simple extension* and proceed to the next round. If $x \in H_i$ then we delete one of the edges in H_i incident with y to create a path Q' and then use the edge (x, y) to concatenate Q, Q' to make P_{j+1}. If $y \in R$ then $P_{j+1} = Q + y$.

If $Q = (v_1, v_2, \dots, v_s) \in \mathcal{Q}$ and $(v_s, v_1) \in A_1$ then we can take any $y \notin V_j$ and with probability at least $1 - (1-q)^s = 1 - O(t^{-\omega(1)})$ find an edge $(y, v_i) \in A_2$. If there is a cycle H_i with $H_i \cap V_j = \emptyset$ then we choose $y \in H_i$ and delete one edge of H_i incident with y to create a path Q' and then we can take $P_{j+1} = (Q', v_i, v_{i-1}, \ldots, v_{i+1})$ and proceed to the next round. Failing this, we choose any $y \in R \setminus V_j$ and let $P_{j+1} = (y, v_i, v_{i-1}, \ldots, v_{i+1})$ and proceed to the next round. Note that this is the first time we will have examined the A_2 edges incident with y. We call this a *cycle extension*.

Suppose now that $Q = (v_1, v_2, \ldots, v_s) \in \mathcal{Q}$ and $(v_s, v_i) \in A_1$ where $1 < i < s-1$. The path $Q' = (v_1, \ldots, v_i, v_s, v_{s-1}, \ldots, v_{i+1})$ is said to be obtained by a rotation. v_1 is the *fixed* endpoint. We partition $\mathcal{Q} = \mathcal{Q}_0 \cup \mathcal{Q}_1 \cup \cdots \cup \mathcal{Q}_{k_0}, k_0 = \log t$ where $\mathcal{Q}_0 = \{P_j\}$ and $\mathcal{Q}_i$ is the set of paths that are obtainable from P_j by exactly i rotations with fixed endpoint v_1. We let N_i denote the set of endpoints of the paths in $\mathcal{Q}_i$, other than v_1, and let $\nu_i = |N_i|$ and let $N_{\mathcal{Q}} = \bigcup_i N_i$. We will prove that q.s.

$$|\nu_i| \leq \frac{1}{100q} \text{ implies that } |\nu_{i+1}| \geq \frac{|\nu_i| t^d q}{300}. \tag{37}$$

It follows from this that q.s. we either end the round through a simple or cycle extension or arrive at a point where the paths in $\mathcal{Q}$ have $\Omega(t^d)$ distinct endpoints. We can take an arbitrary $y \notin V_j$ and find an A_2 neighbor of y among $N_{\mathcal{Q}}$. The probability we cannot find a neighbor is at most $(1-q)^{\Omega(t^d)} = O(t^{-\omega(1)})$. Once we prove (37) we will have shown that we can create a Hamilton path through $\mathcal{Y}_{t,p}^d$ from $H_1, H_2, \ldots, H_\ell, R$ at an extra cost of

$$O(d^{1/2} t \times \log t \times (\ell + t^{d-1}/\log t) = O(t^d).$$

Explanation: Each edge added because of a rotation or through extending the current path costs at most $d^{1/2}t$. The $\log t$ factor comes from the fact that each path is obtained by at most k_0 rotations. We extend into another cycle at most $\ell - 1$ times and into R at most $t^{d-1}/\log t$ times.

We will not have used any A_3 edges to do this.

Proof of (37): We first prove that in the graph induced by A_1 we q.s. have

$$|S| \leq \frac{1}{100q} \text{ implies that } |N_{A_1}(S)| \geq \frac{|S| t^d q}{100}. \tag{38}$$

Here $N_{A_1}(S)$ is the set of vertices not in S that have at least one A_1-neighbor in S.

Indeed, if $s_0 = \frac{1}{100q} = o(n = t^d)$ then

$$\begin{aligned}
\mathbf{Pr}\left(\exists S : |N_{A_1}(S)| < \frac{|S| t^d q}{100}\right) &\leq \sum_{s=1}^{s_0} \binom{t^d}{s} \mathbf{Pr}\left(Bin(t^d - s, 1 - (1-q)^s) \leq \frac{s t^d q}{100}\right) \\
&\leq \sum_{s=1}^{s_0} \binom{t^d}{s} \mathbf{Pr}\left(Bin\left(t^d - s, \frac{sq}{2}\right) \leq \frac{s t^d q}{100}\right) \\
&\leq \sum_{s=1}^{s_0} \left(\frac{t^d e}{s} \cdot e^{-\Omega(t^d q)}\right)^s \\
&= O(t^{-\omega(1)}).
\end{aligned}$$

Now (37) holds for $i = 0$ because q.s. each vertex in $\mathcal{Y}_{t,p}^d$ is incident with at least $t^d q/2$ A_1 edges. Given (38) for $i = 0, 1, \ldots, i-1$ we see that $\nu_1 + \cdots + \nu_{i-1} = o(\nu_i)$. In which case (38) implies that

$$\nu_{i+1} \geq \frac{|N_{A_1}(N_i)| - (\nu_0 + \cdots + \nu_{i-1})}{2} \geq \frac{t^d q \nu_i}{200 + o(1)}$$

completing an inductive proof of (37).

Let P^* be the Hamilton path created above. We now use rotations with v_1 fixed via the edges A_2 to create $\Omega(t^d)$ Hamilton paths with distinct endpoints. We then see that q.s. one of these endpoints is an A_2-neighbor of v_1 and so we get a tour at an additional cost of $O(d^{1/2}t)$.

This completes the proof of Theorem 3.12. □

The upper bound in Theorem 1.2 follows as before by (i) replacing $\mathcal{Y}_{t,p}^d$ by $\mathcal{X}_{n,p}^d$, allowable because our upper bound holds q.s. and $\mathbf{Pr}(|\mathcal{Y}_{t,p}^d| = t^d) = \Omega(t^{-d/2})$ and then (ii) scaling by $n^{-1/d}$ so that we have points in $[0,1]^d$.

4 An algorithm

To find an approximation to a minimum length tour in $\mathcal{X}_{n,p}$, we can use a simple version of Karp's algorithm [13]. We let $m = (n/K\nu_d \log n)^{1/d}$ for some constant $K > 0$ and partition $[0,1]^d$ into m^d subcubes of side $1/m$, as in Lemma 3.8 . The number of points in each subsquare is distributed as the binomial $Bin(n,q)$ where $q = K \log n/n$ and so we have a.a.s. that every subsquare has $K \log n \pm \log n$, assuming K is large enough. The probability that there is no Hamilton cycle in S_α is $O(e^{-Knqp/2})$ and so a.a.s. every subsquare induces a Hamiltonian subgraph. Using the dynamic programming algorithm of Held and Karp [10] we solve the TSP in each subsquare in time $O(\sigma^2 2^\sigma) \leq n^K$, where $\sigma = \sigma_\alpha = |S_\alpha \cap \mathcal{X}_{n,p}|$. Having done this, we can with probability of failure bounded by $m^2(1-p^2)^{(K\log n)^2}$ patch all of these cycles into a tour at an extra $O(m^{d-1}) = o(n^{\frac{d-1}{d}})$ cost. The running time of this step is $O(m^d \log^2 n)$ and so the algorithm is polynomial time overall. The cost of the tour is bounded q.s. as in Lemma 3.8. This completes the proof of Theorem 1.4.

5 Further questions

Theorem 1.1 shows that there is a definite qualitative change in the diameter of $\mathcal{X}_{n,p}$ at around $p = \frac{\log^d n}{n}$, but our methods leave a $(\log\log n)^{2d}$ size gap for the thresholds.

1. What is the precise threshold for there to be distances in $\mathcal{X}_{n,p}$ which tend to ∞? What is the precise threshold for distance in $\mathcal{X}_{n,p}$ to be arbitrarily close to Euclidean distance? What is the behavior of the intermediate regime?

One could also analyze the geometry of the geodesics in $\mathcal{X}_{n,p}$ (Figure 1). For example:

2. Let ℓ be the length of a random edge on the geodesic between fixed points at at constant distance in $\mathcal{X}_{n,p}$. What is the distribution of ℓ?

Improving Theorem 1.2 to give an asymptotic formula for $T(\mathcal{X}_{n,p})$ is another obvious target. It may seem unreasonable to claim such a formula for all (say, decreasing) functions p; in particular, in this case, the constant in the asymptotic formula would necessarily be universal. The following, however, seems reasonable:

Conjecture 5.1. *If $p = \frac{1}{n^\alpha}$ for some constant $0 < \alpha < 1$ then there exists a constant $\beta_{\alpha,d}$ such that a.a.s. $T(\mathcal{X}_{n,p}) \approx \beta_{\alpha,d}\frac{n^{\frac{d-1}{d}}}{p^{1/d}}$.*

We note that $T(\mathcal{X}_{n,1})$ is known to be remarkably well-concentrated around its mean; see, for example, the sharp deviation result of Rhee and Talagrand [19].

3. How concentrated is the random variable $T(\mathcal{X}_{n,p})$?

The case of where $p = o(1)$ may be particularly interesting.

Even for the case $p = 1$ covered by the BHH theorem, the constant $\beta_{1,d}$ $(d \geq 2)$ from Theorem 3.1 is not known. Unlike the case of $p = 1$, the 1-dimensional case is not trivial for our model. In particular, we have proved Theorems 1.3 and 1.2 only for $d \geq 2$. We have ignored the case $d = 1$ not because we consider the technical problems insurmountable, but because we hope that it may be possible to prove a stronger result for $d = 1$, at least for the case of constant p.

4. Determine an explicit constant $\beta_{p,1}$ as a function of (constant) p such that for $d = 1$,

$$\lim_{n\to\infty} T(\mathcal{X}_{n,p}) = \beta_{p,1}.$$

Our basic motivation has been to understand the constraint imposed on travel among random points by the restriction set of traversable edges which is chosen randomly independently of the geometry of the underlying point-set. While the Erdős-Rényi-Gilbert model is the prototypical example of a random graph, other models such as the Barabási-Albert preferential attachment graph have received wide attention in recent years, due to properties (in particular, the distribution of degrees) they share with real-world networks. In particular, if the random graph one is traveling within is the flight-route map for an airline, the following questions may be the most relevant:

5. If the preferential attachment graph is embedded randomly in the unit square (hypercube), what is the expected diameter? What is the expected size of a minimum-length spanning tree?

Similarly, one could examine a combination of geometry and randomness in determining connections in the embedded graph. Our methods already give something in this direction. In particular, we can define $\mathcal{X}_{n,p,r}$ as the intersection of the graphs $\mathcal{X}_{n,p}$ with the random geometric graph on the vertex set $\mathcal{X}_n$, where a pair of points are joined by an edge if they are at distance $\leq r$. Following our proof of Theorem 1.3, one sees that we find that

Theorem 5.2. *If $d \geq 2$, $p > 0$ is constant, and $r = r(n) \geq n^{\varepsilon - 1/d}$ for some $\varepsilon > 0$, then*

$$T(\mathcal{X}_{n,p,r}) \approx \beta_{p,d} n^{\frac{d-1}{d}} \qquad a.a.s.$$

This is because we have shown that a.a.s. we can find a near optimum tour that only uses edges much smaller than the radius r.

Of course, the ideas behind Question 5 and Theorem 5.2 could be considered together; note that Flaxman, Frieze and Vera [6] considered a geometric version of a preferential attachment graph.

The proof of Theorem 1.4 is relatively painless. We are reminded that Arora [1] and Mitchell [16] have described more sophisticated polynomial time algorithms that are asymptotically optimal even with the worst-case placing of the points. It would be interesting to see whether these algorithms can handle the random loss of edges.

6. Do the methods of Arora and Mitchell allow efficient approximation of the tour length through $\mathcal{X}_{n,p}$, when the embedding $\mathcal{X}_n$ is *arbitrary*?

References

[1] S. Arora, Polynomial time approximation schemes for Euclidean Traveling Salesman and other geometric problems, *Journal of the Association for Computing Machinery* 45 (1998) 753-782.

[2] J. Beardwood, J. H. Halton and J. M. Hammersley, The shortest path through many points, *Mathematical Proceedings of the Cambridge Philosophical Society* 55 (1959) 299-327.

[3] B. Bollobás, Random Graphs, First Edition, Academic Press, London 1985, Second Edition, Cambridge University Press, 2001.

[4] B. Bollobás, T. Fenner and A.M. Frieze, An algorithm for finding Hamilton paths and cycles in random graphs, *Combinatorica* 7 (1987) 327-341.

[5] D. Fernholz and V. Ramachandran, The diameter of sparse random graphs, *Random Structures and Algorithms* 31 (2007) 482-516.

[6] A. Flaxman, A.M. Frieze and J. Vera, A Geometric Preferential Attachment Model of Networks *Internet Mathematics* 3 (2007) 187-205.

[7] A.M. Frieze, On large matchings and cycles in sparse random graphs, *Discrete Mathematics* 59 (1986) 243-256.

[8] A.M. Frieze and M. Karoński, Introduction to Random Graphs, Cambridge University Press, 2015.

[9] Y. Gurevich and S. Shelah, Expected computation time for Hamiltonian path problem, *SIAM Journal on Computing* 16 (1987) 486-502.

[10] M. Held and R.M. Karp, A dynamic programming approach to sequencing problems, *SIAM Journal on Applied Mathematics* 10 (1962) 196-210.

[11] W. Hoeffding, Probability inequalities for sums of bounded random variables, *Journal of the American Statistical Association* 58 (1963) 13-30.

[12] A. Johansson, J. Kahn and V. Vu, Factors in Random Graphs, *Random Structures and Algorithms* 33 (2008) 1-28.

[13] R.M. Karp, Probabilistic Analysis of Partitioning Algorithms for the Traveling-Salesman Problem in the Plane, *Mathematics of Operations Research* 2 (1977) 209-244.

[14] A. Mehrabian, A Randomly Embedded Random Graph is Not a Spanner, In Proceedings of the 23rd Canadian Conference on Computational Geometry (CCCG 2011) (2011) 373-374.

[15] A. Mehrabian and N. Wormald, On the Stretch Factor of Randomly Embedded Random Graphs, to appear.

[16] J. Mitchell, Guillotine Subdivisions Approximate Polygonal Subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems, *SIAM Journal on Computing* 28 (1999) 1298-1309.

[17] M. Penrose, Random Geometric Graphs, Oxford University Press, 2003.

[18] L. Pósa, Hamiltonian circuits in random graphs, *Discrete Mathematics* 14 (1976) 359-364.

[19] W. Rhee and M. Talagrand, A Sharp Deviation Inequality for the Stochastic Traveling Salesman Problem, in *The Annals of Probability* 17 (1989) 1–8.

[20] O. Riordan and N. Wormald, The diameter of sparse random graphs, *Combinatorics, Probability and Computing* 19 (2010) 835-926.

[21] J. Michael Steele, Subadditive Euclidean functionals and nonlinear growth in geometric probability, *The Annals of Probability* 9 (1981) 365-376.

[22] J. Yukich, Probability Theory of Classical Euclidean Optimization Problems, Springer, 1991.

Appendix: Proof of Observation 3.2

Assume without loss of generality that we have scaled so that $\mu = 1$. We write, for $\lambda > 0$ such that $e^\lambda < 1/\rho$,

$$\mathbf{E}(Y^2 e^{\lambda Y}) = \sum_{k=0}^{\infty} k^2 e^{\lambda k}\, \mathbf{Pr}(Y = k) \le C \sum_{k=0}^{\infty} k^2 (\rho e^\lambda)^k \le \frac{3C}{(1-\rho e^\lambda)^3}.$$

Now $e^x \le 1 + x + x^2 e^x$ and so, using the above, we have

$$\mathbf{E}(e^{\lambda Y}) \le 1 + \lambda + \lambda^2 \left(1 + \frac{3C}{(1-\rho e^\lambda)^3}\right).$$

So, if $Z = Y_1 + Y_2 + \cdots + Y_n$ where $Y_1, Y_2, \ldots, Y_n$ are independent copies of Y,

$$\begin{aligned}\mathbf{Pr}(Z \ge n + \delta n) &\le e^{-\lambda(1+\delta)n}\, \mathbf{E}(e^{\lambda Y})^n \\ &\le e^{-\lambda(1+\delta)n} \exp\left\{\left(\lambda + \lambda^2 \left(1 + \frac{3C}{(1-\rho e^\lambda)^3}\right)\right) n\right\} \\ &\le e^{-\lambda\delta n} \exp\left\{\lambda^2(1 + 3C\varepsilon^{-3})\right\}\end{aligned}$$

assuming that

$$e^\lambda \le (1-\varepsilon)/\rho. \qquad (39)$$

Now choose $\lambda = \delta/(1 + 3C\varepsilon^{-3})$ and $\varepsilon = \varepsilon(\delta)$ such that (39) holds. Then

$$\mathbf{Pr}\left(Z \ge n + \delta n)\right) \le \exp\left\{-\frac{\delta^2 n}{2(1 + 2\varepsilon^{-3})}\right\}.$$

To bound $\mathbf{Pr}(Z \leq n - \delta n)$ we use

$$\begin{aligned}
\mathbf{Pr}(Z \leq n - \delta n) &\leq e^{\lambda(1-\delta)n}\,\mathbf{E}(e^{-\lambda Y})^n \\
&\leq e^{\lambda(1-\delta)n} \exp\left\{\left(-\lambda + \lambda^2\left(1 + \frac{3C}{(1-\rho e^{\lambda})^3}\right)\right) n\right\} \\
&\leq e^{-\lambda\delta n} \exp\left\{\lambda^2(1 + 3C\varepsilon^{-3})\right\},
\end{aligned}$$

and we can proceed as before.