Disjoint Hamilton cycles in the random geometric graph

Tobias Müller*
Centrum voor Wiskunde en Informatica
P.O. Box 94079
1090 GB Amsterdam, The Netherlands
tobias@cwi.nl

Xavier Pérez-Giménez†
Department of Combinatorics and Optimization
University of Waterloo
Waterloo ON, Canada
xperez@uwaterloo.ca

Nicholas Wormald‡
Department of Combinatorics and Optimization
University of Waterloo
Waterloo ON, Canada
nwormald@uwaterloo.ca

Abstract

We consider the standard random geometric graph process in which n vertices are placed at random on the unit square and edges are sequentially added in increasing order of edge-length. For fixed $k \geq 1$, we prove that the first edge in the process that creates a k-connected graph coincides a.a.s. with the first edge that causes the graph to contain $k/2$ pairwise edge-disjoint Hamilton cycles (for even k), or $(k-1)/2$ Hamilton cycles plus one perfect matching, all of them pairwise edge-disjoint (for odd k). This proves and extends a conjecture of Krivelevich and Müller. In the special when case $k = 2$, our result says that the first edge that makes the random geometric graph Hamiltonian is a.a.s. exactly the same one that gives 2-connectivity, which answers a question of Penrose. We prove our results with lengths measured using the ℓ_p norm for any $p > 1$, and we also extend our result to higher dimensions.

1 Introduction

Many authors have studied the evolution of the random geometric graph on n labelled vertices placed independently and *uniformly at random* (u.a.r.) on the unit square $[0,1]^2$, in which edges are added in increasing order of length (see e.g. [7]). Penrose [6] proved that the first added edge that makes the graph have minimum degree k is *asymptotically almost surely* (a.a.s., i.e. with probability tending to 1 as $n \to \infty$) the first one that makes it k-connected. [**TM: rewrote**] Here the edge-lengths could be measured not only by the Euclidean norm, but also using other norms. For technical reasons, Penrose required the edge-lenghts to be measured using an ℓ_p norm with $p > 1$ in this last result.

Penrose also asked whether, in the evolution of the random geometric graph, 2-connectivity occurs a.a.s. precisely when the first Hamilton cycle is created. As a first step towards answering Penrose's question, Díaz, Mitsche and the second author showed in [3] that the property

*Research partially supported by a VENI grant from Netherlands Organisation for Scientific Research (NWO).
†Partially supported by the Province of Ontario under the Post-Doctoral Fellowship (PDF) Program.
‡Supported by the Canada Research Chairs Program and NSERC.

of being Hamiltonian has a sharp threshold at $r \sim \sqrt{\log n/(\theta n)}$ (here θ denotes the area of the unit ball with respect to the norm used), which coincides asymptotically with the threshold for k-connectivity for any constant k. On the other hand, a more general result is already known for the evolution of the random graph G on n labelled vertices, in which edges are added one by one. Bollobás and Frieze showed in [2] that a.a.s. as soon as G has minimum degree k, it also contains $\lfloor k/2 \rfloor$ edge-disjoint Hamilton cycles plus an additional edge-disjoint perfect matching if k is odd, where k is any constant positive integer. The main result in the present paper is that the analogue of Bollobás and Frieze's result holds for the random geometric graph. That is, we show that, in the evolution of the random geometric graph, a.a.s. as soon as the graph becomes k-connected, it immediately contains $\lfloor k/2 \rfloor$ edge-disjoint Hamilton cycles plus one additional perfect matching if k is odd. This result applies to the geometric graph in a unit hypercube of $d \geq 2$ dimensions, and for the ℓ_p norm, $1 < p \leq \infty$. This proves and extends a conjecture of Krivelevich and the first author [4], who conjectured the result for even values k. The special case $k = 2$ of our result answers Penrose's question.

Three independent but similar proofs for Penrose's question appeared in preprints by Balogh, Bollobás and Walters [1], by Krivelevich and Müller [4] and by Pérez-Giménez and Wormald [8]. The present paper presents the proof in [8], which is only for dimension 2 but does cover arbitrary k, and includes the extension to arbitrary d making use of a result in [4].

Ensuring that edge-disjoint subgraphs exist requires various modifications of the argument for a single cycle. An important difficulty is that the cycles need to visit, edge-disjointly, those unusual sets of vertices that have relatively few common neighbours. Note that k-connectedness assures us that $\lfloor k/2 \rfloor$ edge-disjoint paths can pass through a single vertex, but can those paths be extended far if the vertex lies in a clique of $j < k$ vertices in a k-connected graph, when the j vertices have a 'small' common neighbourhood? We call such sets of vertices 'dangerous,' and they constitute the main obstacle in our construction. These considerations gave rise to questions about nonrandom graphs that are not apparently answered by known results. In this paper we give two alternative treatments of the small vertex cuts. The first treatment proves a result on packing linear forests in graphs together with a (weak) probabilistic characterisation (Lemma 4) of what the small vertex cuts can look like. This shows that certain local structures do not immediately exclude the existence of $\lfloor k/2 \rfloor$ edge-disjoint Hamilton cycles in a k-connected graph. This may be of independent interest, and also brings up the question of what is the most likely reason that a k-connected random geometric graph does *not* have the required set of Hamilton cycles. The second, shorter treatment of small vertex cuts makes do with a much weaker packing result by excluding more situations probabilistically (see Lemma 9).

Let $\boldsymbol{X} = (X_1, \ldots, X_n)$ be a random vector, where each X_i is a point in $[0,1]^2$ chosen independently with uniform distribution. Given $\boldsymbol{X}$ and a radius $r = r(n) \geq 0$, we define the *random geometric graph* $\mathscr{G}(\boldsymbol{X}; r)$ as follows: the vertex set of $\mathscr{G}(\boldsymbol{X}; r)$ is $\{1, \ldots, n\}$ and there is an edge joining i and j whenever $\|X_i - X_j\|_p \leq r$. Here $\|\cdot\|_p$ denotes the standard ℓ_p norm, for some fixed $1 < p \leq \infty$. Unless otherwise stated, all distances in $[0,1]^2$ are measured according to the ℓ_p norm (i.e. $d(X,Y) = \|X - Y\|_p$). Let θ be the area of the unit ℓ_p-ball (e.g. $\theta = \pi$ for $p = 2$, and $2 \leq \theta \leq 4$ for all $1 \leq p \leq \infty$).

A continuous-time random graph process $\big(\mathscr{G}(\boldsymbol{X}; r)\big)_{0 \leq r < \infty}$ is defined in a natural way, by first choosing the random set of points $\boldsymbol{X}$ and then adding edges one by one as we increase the radius r from 0 to ∞.

Theorem 1. *Consider the random graph process* $\big(\mathscr{G}(\boldsymbol{X}; r)\big)_{0 \leq r < \infty}$ *for any* ℓ_p*-normed metric on* $[0,1]^2$*,* $1 < p \leq \infty$*, and let* k *be a fixed positive integer.*

(i) For even $k \geq 2$, a.a.s. the minimum radius r at which the graph $\mathscr{G}(\boldsymbol{X}; r)$ is k-connected is equal to the minimum radius at which it has $k/2$ edge-disjoint Hamilton cycles.

(ii) For odd $k \geq 1$, a.a.s. the minimum radius r at which the graph $\mathscr{G}(\boldsymbol{X}; r)$ is k-connected is equal to the minimum radius at which it has $(k-1)/2$ Hamilton cycles and one perfect matching, all of them pairwise edge-disjoint. (Here asymptotics are restricted to even n.)

[TM: added next sentence] The reason that we restrict ourselves to the ℓ_p norm with $p > 1$ to measure the edge-lengths in Theorem 1 (as opposed to a completely arbitrary norm), is that this restriction is imposed by the results of Penrose that we invoke in the proof of Theorem 1.

Our proof of Theorem 1 has two versions. In the first, we prove a deterministic result about packing paths in graphs. A *linear forest* is a forest all of whose components are paths. We use $d_G(v)$ and $N_G(v)$ to denote the degree and set of neighbours (respectively) of a vertex v in a graph G.

Lemma 2. *Assume $k \geq 1$, $1 \leq j \leq k$ and $\ell = k - j + 1$. Let G be a graph with vertex set $J \cup B$ with $|J| = j$ consisting of a clique on vertex set J together with a bipartite graph H with parts J and B, such that*

(i) $d_H(v) \geq \ell$ for each $v \in J$, and there exists a special vertex we call the apex *which has degree at least $\ell + 1$ in H;*

(ii) for each pair of distinct vertices $v, v' \in J$, $|N_H(\{v, v'\}) \setminus \{v, v'\}| \geq \ell + 1$.

Then G contains $\lfloor k/2 \rfloor$ pairwise edge-disjoint linear forests, and additionally, if k is odd, a matching, such that each edge of the clique on vertex set J is contained in one of the forests or in the matching, and each vertex in J has degree 2 in each forest and (for odd k) degree 1 in the matching.

This result is combined with a probabilistic lemma about the likely neighbourhoods of dangerous sets of vertices (see Lemma 4 in Section 2), which guarantees the existence of an apex vertex. We conjecture that the assumption about the apex vertex can be removed from Lemma 2.

Conjecture 1. *Lemma 2 is valid with (i) replaced simply by the condition*

(i) $d_H(v) \geq \ell$ for each $v \in J$.

In the second version of the proof of the main theorem, we use a stronger probabilistic lemma. This allows us to assume a much stronger property of the neighbourhoods of the dangerous sets of vertices described above, which greatly simplifies the packing problem and does not require Lemma 2. The first proof is considerably longer as a whole but does address the question of what local conditions in k-connected graphs can prevent the existence of edge-disjoint Hamilton cycles. We find this interesting in the context of investigating random geometric graphs that are k-connected. Lemma 2 has an impact on such an investigation, and Conjecture 1, if true, would have an even stronger impact, as discussed at the end of Section 4.

The next section contains the basic geometric definitions and probabilistic statements required in the argument, including proofs that several properties hold a.a.s. Next, we prove Lemma 2 in Section 3. Then, in Section 4, we prove the main theorem, by supplying the required construction of Hamilton cycles (and perfect matching) in the random geometric graph deterministically, assuming the properties that were shown to hold a.a.s. Additionally, Section 5 gives an alternative version of the proof not requiring Lemma 2, and finally, in Section 6, we extend the argument to general dimension.

2 Asymptotically almost sure properties

Let $k \geq 1$ be a fixed integer, and define $m = 2k - 3$ if $k \geq 2$ and $m = 0$ if $k = 1$. We state a result which is a consequence of Theorem 8.4 in [7].

Proposition 3. *In the random process* $\big(\mathscr{G}(\boldsymbol{X}; r)\big)_{0 \leq r < \infty}$*, let* r_k *be the smallest* r *such that* $\mathscr{G}(\boldsymbol{X}; r)$ *is* k*-connected. Then,*

$$\theta n {r_k}^2 - \log n - m \log\log n \tag{1}$$

is bounded in probability.

Here, we define some properties of the random geometric graph that hold a.a.s. and that will turn out to be sufficient for our construction of disjoint Hamilton cycles. In view of Proposition 3, we shall mainly focus our analysis to r satisfying $\theta n r^2 = \log n + m \log\log n + O(1)$ or sometimes just $\theta n r^2 \sim \log n$.

Henceforth we assume that the points in $\boldsymbol{X}$ are in general position—i.e. they are are all different, no three of them are collinear, and all distances between pairs of points are strictly different—since this holds with probability 1.

Lemma 4. *For any small enough constant* $\eta > 0$ *and for any* r *such that* $\theta n r^2 = \log n + m \log\log n + O(1)$*, the random geometric graph* $\mathscr{G}(\boldsymbol{X}; r)$ *a.a.s. satisfies the following property. Every set* J *of vertices of size* $2 \leq |J| \leq k$ *in which each vertex has degree at least* k *and such that* $\max_{u,v \in J}\{d(X_u, X_v)\} \leq \eta r$ *contains some vertex of degree at least* $k + 1$*.*

Proof. Suppose that some set $J \subset \{1, \dots, n\}$ with $|J| = j$ causes the property to fail. In particular, the Euclidean diameter of $\{X_v \ : \ v \in J\}$ is at most $\sqrt{2}\eta r$, and all vertices in J must have exactly k neighbours. Let v_1 and v_2 be two vertices of J such that X_{v_1} and X_{v_2} realise such Euclidean diameter, and assume w.l.o.g. that X_{v_1} is closer to the boundary of $[0, 1]^2$ than X_{v_2} is. By construction, v_1 has exactly $k + 1 - j$ neighbours not in J, and the number s of neighbours of v_2 which are not neighbours of v_1 satisfies $s \leq k - j + 1$. Thus, we have an example of one of the following configurations. Given a fixed nonnegative integer s, a *bad* configuration is an ordered tuple J of $k + s + 1$ vertices $v_1, \dots, v_{k+s+1}$ in $\mathscr{G}(\boldsymbol{X}; r)$ with the following properties: vertices $v_2, \dots, v_{k+1}$ are the only k neighbours of v_1 in $\mathscr{G}(\boldsymbol{X}; r)$; X_{v_1} and X_{v_2} are at Euclidean distance $\rho \leq \sqrt{2}\eta r$; X_{v_1} is closer to the boundary of $[0, 1]^2$ than X_{v_2} is; all $X_{v_2}, \dots, X_{v_j}$ lie in the circle of radius ρ centred at X_{v_1}; vertex v_2 has exactly s neighbours that are not neighbours of v_1, namely $v_{k+2}, \dots, v_{k+s+1}$. To prove the statement it suffices to show that, for any fixed s, a.a.s. bad configurations do not occur.

We distinguish three types of bad configurations according to the position of X_{v_1} in $[0, 1]^2$: call *corner* bad configurations to the ones in which X_{v_1} is at distance at most r from two of the four sides of $[0, 1]^2$; let *side* bad configurations be the ones in which X_{v_1} is at distance at most r from exactly one of the four sides of $[0, 1]^2$; all other bad configurations are referred to as *interior* bad configurations. Let T_1, T_2 and T_3 denote respectively the number of corner, side and interior bad configurations. We only have to show that $\mathbf{E}(T_1 + T_2 + T_3) = o(1)$.

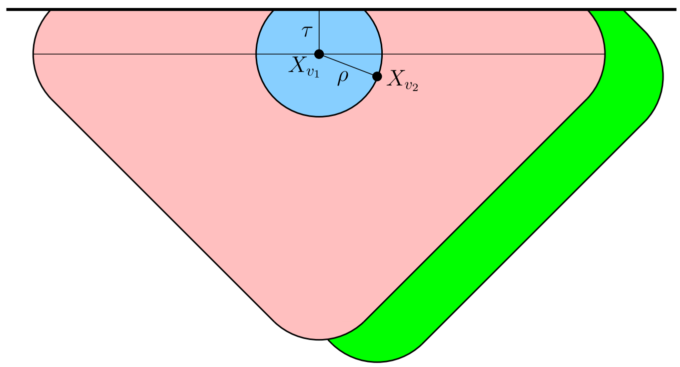

Figure 1: Visual description of a bad configuration. Points $X_{v_3}, \dots, X_{v_j}$ must lie in the blue area. Points $X_{v_{j+1}}, \dots, X_{v_{k+1}}$ must lie in the blue or pink areas. Points $X_{v_{k+2}}, \dots, X_{v_{k+s+1}}$ must lie in the green area.

First we study in detail the side bad configurations. Assume we are given the position of X_{v_1} and $\rho \leq \sqrt{2}\eta r$. Let $\tau \leq r$ be the distance between X_{v_1} and the closest side of $[0,1]^2$. The probability that vertices $v_1, \dots, v_{k+s+1}$ form a side bad configuration can be easily bounded in terms of ρ and τ by estimating the area of the regions where each vertex must lie (see Figure 1). The probability that X_{v_i} lies in the required circle for $2 < i \leq j$ is at most $(\pi\rho^2)^{j-2} = O(\rho^{2j-4})$. The probability that $d(v_1, v_i) \leq r$ for $j+1 \leq i \leq k+1$ is at most $(\theta r^2)^{k+1-j} = O(r^{2(k+1-j)})$. It can be seen that, provided η was chosen sufficiently small, we have

$$\frac{1}{2}\rho r \leq \text{area}\left([0,1]^2 \cap B_p(X_{v_2}, r) \setminus B_p(X_{v_1}, r)\right) \leq 3\rho r,$$

where $B_p(x,s) := \{y \in \mathbb{R}^2 : \|x-y\|_p < s\}$ denotes the l_p-ball around x of radius s. Thus the probability that vertices $v_{k+2}, \dots, v_{k+s+1}$ are neighbours of v_2 but not v_1 is at most $(3\rho r)^s$. The probability that the remaining $n-k-s-1$ vertices are not neighbours of v_1 or v_2 is at most

$$(1-\theta r^2/2 - \tau r - \frac{1}{2}\rho r)^{n-k-1} \leq \exp(-(n-k-1)(\theta r^2/2 + \tau r + \frac{1}{2}\rho r)) = O\left(\frac{e^{-\tau r n - \frac{1}{2}\rho r n}}{n^{1/2}\log^{m/2} n}\right).$$

The probability density function of the distance between X_{v_1} and the closest side of $[0,1]^2$ is at most 4. The probability density function of the Euclidean distance between X_{v_1} and X_{v_2} is at most $2\pi\rho$. Putting all together and multiplying by the number of choices of the vertices $v_1, \dots, v_{k+s+1}$,

$$\begin{aligned}
\mathbf{E}T_2 &= O(n^{k+s+1}) \int_0^r \int_0^{\sqrt{2}\eta r} \rho^{2j-4} r^{2(k+1-j)} (\rho r)^s \frac{e^{-\tau r n - \frac{1}{2}\rho r n}}{n^{1/2}\log^{m/2} n} \rho d\rho d\tau \\
&= O(1) \frac{(nr^2)^{k+3/2-2j}}{\log^{m/2} n} \int_0^{r^2 n} e^{-y} dy \int_0^{\sqrt{2}\eta\frac{1}{2}r^2 n} x^{2j+s-3} e^{-x} dx \\
&= O(\log^{3-2j+(2k-3-m)/2} n) = o(1),
\end{aligned}$$

since $j \geq 2$ and $m \geq 2k-3$. Here we have used the substitutions $y = \tau r n$ and $x = \frac{1}{2}\rho r n$ to get the second line.

The analysis of interior bad configurations is fairly analogous and in fact simpler, since we need not integrate with respect to τ. We omit the details and simply state that

$$\begin{aligned}\mathbf{E}T_1 &= O(n^{k+s+1})\int_0^{\sqrt{2}\eta r} \rho^{2j-4} r^{2(k+1-j)}(\rho r)^s \frac{e^{-\frac{1}{2}\rho r n}}{n\log^m n}\rho d\rho \\ &= O(1)\frac{(nr^2)^{k+2-2j}}{\log^m n}\int_0^{\sqrt{2}\eta\frac{1}{2}r^2 n} x^{2j+s-3}e^{-x}dx \\ &= O(\log^{3-2j+(k-1-m)} n) = o(1),\end{aligned}$$

since $j \geq 2$ and $m \geq k-1$. Finally,

$$\mathbf{E}T_3 = O(n^{k+s+1})\frac{r^{2(k+s+1)}}{n^{1/4}\log^{m/4} n} = O\left(\frac{\log^{k+s+1-m/4} n}{n^{1/4}}\right) = o(1). \qquad \square$$

For the following definitions, we fix $\delta > 0$ to be a small enough constant and assume $r \to 0$. We tessellate $[0,1]^2$ into square *cells* of side $\delta' r = \lceil (\delta r)^{-1} \rceil^{-1}$. (Note that δ' is not constant, but $\delta' \leq \delta$ and $\delta' \to \delta$). Let $\mathcal{C}$ be the set of cells, and let $\mathscr{G}_{\mathcal{C}}$ be an auxiliary graph with vertex set $\mathcal{C}$ and with one edge connecting each pair of cells c_1 and c_2 iff all points in c_1 have distance at most r from all points in c_2. Note that we shall use the term *adjacent cells* to refer to cells which are adjacent vertices of the graph of cells $\mathscr{G}_{\mathcal{C}}$, while cells sharing a side boundary will be described as being *topologically adjacent*. Let Δ be the maximum degree of $\mathscr{G}_{\mathcal{C}}$. By construction, Δ is a constant only depending on δ and the chosen ℓ_p norm.

We may assume that each point in $\boldsymbol{X}$ lies strictly in the interior of a cell in the tesselation, since this happens with probability 1. Let M be a large enough but constant positive integer (its choice will only depend on Δ, thus on δ, and also on k and ℓ_p). A cell in $\mathcal{C}$ is *dense* if it contains at least M points of the random set $\boldsymbol{X}$, *sparse* if it contains at least one, but less than M, points in $\boldsymbol{X}$, and *empty* if it has no points in $\boldsymbol{X}$. Let $\mathcal{D} \subseteq \mathcal{C}$ be the set of dense cells. Note that $\mathcal{D} \neq \emptyset$, since the total number of cells is $|\mathcal{C}| = \Theta(n \log n)$, so at least one must contain $\Omega(\log n)$ points in $\boldsymbol{X}$.

A set of cells is said to be *connected* if it induces a connected subgraph of $\mathscr{G}_{\mathcal{C}}$. (For δ small enough, this includes the situation where the union of cells is topologically connected.) The *area* of a set of cells is simply the area of the corresponding union of cells. A set of cells *touches* one side (or one corner) of $[0,1]^2$ if it contains a cell which has some boundary on that side (or corner) of the unit square.

Lemma 5. *For any constants $\delta > 0$ and $\alpha > 0$ and for any r satisfying $\theta n r^2 \sim \log n$, the following statements hold a.a.s.*

1. *All connected sets of cells of area at least $(1+\alpha)\theta r^2$ contain some dense cell.*
2. *All connected sets of cells of area at least $(1+\alpha)\theta r^2/2$ touching some side of $[0,1]^2$ contain some dense cell.*
3. *All cells contained inside a $5r \times 5r$ square on each corner of $[0,1]^2$ are dense.*

Proof. Recall that the area of each cell is $\delta'^2 r^2$. Then, in order to show the first statement in the lemma, it suffices to consider all connected sets of cells with exactly $s = \lceil (1+\alpha)\theta/\delta'^2 \rceil = \Theta(1)$ cells. Let $\mathcal{S}$ be such a set of cells. The probability that $\mathcal{S}$ has no dense cell is at most

$$\begin{aligned}\sum_{t=0}^{(M-1)s}\binom{n}{t}(s\delta'^2 r^2)^t(1-s\delta'^2 r^2)^{n-t} &= O\left(e^{-(1+\alpha)\theta r^2 n}\right)\sum_{t=0}^{(M-1)s}(r^2 n)^t \\ &= O\left(n^{-(1+\alpha)+o(1)}\log^{(M-1)s} n\right). \qquad (2)\end{aligned}$$

To conclude the first part of the proof, multiply the probability above by the number $\Theta(1/r^2) = \Theta(n/\log n)$ of connected sets of s cells.

By a completely analogous argument, if $\mathcal{S}$ has area only $\lceil(1+\alpha)\theta/\delta'^2\rceil/2$ and touches some side of $[0,1]^2$, the probability that it has no dense cell is $O(n^{-(1+\alpha)/2)+o(1)}\log^{(M-1)s} n$. However, the number of such sets is only $\Theta(\sqrt{n/\log n})$.

Finally, there is a bounded number of cells inside any of the $5r \times 5r$ squares on the corners, and each individual cell is dense with probability $1 - o(1)$. □

A set of cells is *small* if it can be embedded in a 16×16 grid of cells, and it is *large* otherwise. Consider the subgraph $\mathscr{G}_{\mathcal{C}}[\mathcal{D}]$ of $\mathscr{G}_{\mathcal{C}}$ induced by dense cells, and let $\mathcal{D}_0$ be the set of dense cells which are not in small components of $\mathscr{G}_{\mathcal{C}}[\mathcal{D}]$ (we shall see that $\mathcal{D}_0$ forms a unique large component in $\mathscr{G}_{\mathcal{C}}[\mathcal{D}]$). Most of the trouble in our argument comes from cells which are not adjacent to any dense cell in $\mathcal{D}_0$, so let $\mathcal{B} = \mathcal{C} \setminus (\mathcal{D}_0 \cup N(\mathcal{D}_0))$, and call the cells in $\mathcal{B}$ *bad* cells. Also, let us denote components of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$ as *bad* components. Note that by construction all cells in $N(\mathcal{B}) \setminus \mathcal{B}$ must be sparse but adjacent to some cell in $\mathcal{D}_0$, while $\mathcal{B}$ itself may contain both sparse and dense cells.

Lemma 6. *For a small enough constant $\delta > 0$ and for any r satisfying $\theta n r^2 \sim \log n$, the following holds a.a.s.*

1. *All components of $\mathscr{G}_{\mathcal{C}}[\mathcal{D}]$ are small except for one large component formed by precisely the cells in $\mathcal{D}_0$.*
2. *$\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$ has only small components.*

Proof. First, we claim that the following statements are a.a.s. true. (Recall that "connected" is defined in terms of the graph $\mathscr{G}_{\mathcal{C}}$, not topological adjacency.)

1. For any large connected set of cells $\mathcal{S}$ such that $N(\mathcal{S})$ does not touch all four sides of $[0,1]^2$, $N(\mathcal{S}) \setminus \mathcal{S}$ must contain some dense cell.
2. For any pair of connected sets of cells $\mathcal{S}_1$ and $\mathcal{S}_2$ not adjacent to each other (i.e. $\mathcal{S}_2 \cap N(\mathcal{S}_1) = \emptyset$) and such that both $N(\mathcal{S}_1)$ and $N(\mathcal{S}_2)$ touch all four sides of $[0,1]^2$, $N(\mathcal{S}_1) \setminus \mathcal{S}_1$ or $N(\mathcal{S}_2) \setminus \mathcal{S}_2$ must contain some dense cell.

As an immediate consequence of this claim, by considering the maximal connected sets of dense cells, we deduce that $\mathscr{G}_{\mathcal{C}}[\mathcal{D}]$ must have a unique large component, consisting of all cells in $\mathcal{D}_0$ (note that $\mathcal{D}_0 \neq \emptyset$ by statement 3 in Lemma 5). Moreover, $N(\mathcal{D}_0)$ must touch all four sides of $[0,1]^2$. Now suppose that $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$ has some large component $\mathcal{S}$. By definition $N(\mathcal{S}) \setminus \mathcal{S}$ contains only sparse cells. Then, by the first part of the claim, $N(\mathcal{S})$ must touch the four sides of $[0,1]^2$. Hence, we apply the second part of the claim to $\mathcal{S}$ and $\mathcal{D}_0$ to deduce that such large $\mathcal{S}$ cannot exist.

It just remains to prove the initial claim. Let $\mathcal{S}$ be a connected set of cells. Observe that $\bigcup N(\mathcal{S})$ is topologically connected (and in particular $N(\mathcal{S})$ is a connected set of cells), and that the outer boundary γ of $\bigcup N(\mathcal{S})$ is a simple closed polygonal path along the grid lines in $[0,1]^2$ defined by the tessellation. If we remove from γ the segments that coincide with some side of $[0,1]^2$, each connected polygonal path that remains is called a *piece* of γ. Note that $N(\mathcal{S}) \setminus \mathcal{S}$ need not be a connected set of cells. However all cells in $N(\mathcal{S})$ along the same piece of γ must be contained in the same topological component of $\bigcup(N(\mathcal{S}) \setminus \mathcal{S})$, and thus in the same connected component of $\mathscr{G}_{\mathcal{C}}[N(\mathcal{S}) \setminus \mathcal{S}]$.

The argument comprises several cases. For each case, a lower bound on the area of some connected component of $\mathscr{G}_{\mathcal{C}}[N(\mathcal{S}) \setminus \mathcal{S}]$ is given by finding some disjoint subsets of $[0,1]^2$ of large enough area contained in the union of cells in that component. Then, Lemma 5 ensures that $N(\mathcal{S}) \setminus \mathcal{S}$ contain at least one dense cell.

Given a cell c, let $B_{\nearrow}(c)$ be the set of points at distance at most $(1-4\delta')r$ from the top right corner of c and above and to the right of that corner. The sets $B_{\nwarrow}(c)$, $B_{\searrow}(c)$ and $B_{\swarrow}(c)$ are defined analogously replacing (top, above, right) by (top, above, left), (bottom, below, right) and (bottom, below, left) respectively. Note that $B_{\nearrow}(c)$, $B_{\nwarrow}(c)$, $B_{\searrow}(c)$ and $B_{\swarrow}(c)$ are disjoint and contained in $\bigcup(N(c) \setminus \{c\})$.

Case 1. Let $\mathcal{S} \subseteq \mathcal{C}$ be a connected set of cells which is not small and such that $N(\mathcal{S})$ does not touch any side of $[0,1]^2$. Since $\mathcal{S}$ is not small, assume without loss of generality that its vertical extent is greater than $16\delta' r$. Let c_1, c_2, c_3, c_4 be respectively the topmost, bottommost, leftmost and rightmost cells in $\mathcal{S}$ (possibly not all different and not unique). Let $A_{\rightarrow}$ be any rectangle of height $16\delta' r$ and width $(1-20\delta')r$ glued to the right of c_4 and between the top of c_1 and the bottom of c_2. Also choose a similar rectangle $A_{\leftarrow}$ of the same dimensions glued to the left of c_3, and let $A_{\uparrow}$ and $A_{\downarrow}$ be rectangles of height $(1-4\delta')r$ and width $\delta' r$ placed on top of, and below, the cells c_1 and c_2 respectively. By construction, $B_{\nearrow}(c_1)$, $B_{\nwarrow}(c_1)$, $B_{\searrow}(c_2)$, $B_{\swarrow}(c_2)$, $A_{\uparrow}$, $A_{\downarrow}$, $A_{\leftarrow}$ and $A_{\rightarrow}$ are disjoint and are contained in the same topological component of $\bigcup(N(\mathcal{S}) \setminus \mathcal{S})$ (i.e. the one that touches γ), which thus has area at least

$$\theta(1-4\delta')^2 r^2 + 2\delta' r(1-4\delta')r + 32\delta' r(1-20\delta')r \geq \theta r^2(1+\delta'/3).$$

Hence, by Lemma 5, $N(\mathcal{S}) \setminus \mathcal{S}$ must contain some dense cell.

Case 2. Let $\mathcal{S} \subseteq \mathcal{C}$ be a connected set of cells which is not small and such that $N(\mathcal{S})$ touches only one side of $[0,1]^2$ (assume it is the bottom side). This is very similar to Case 1, so we just sketch the main differences in the argument.

If the vertical extent of $\mathcal{S}$ is greater than $16\delta' r$, then proceed as in Case 1 but only consider the sets $B_{\nearrow}(c_1)$, $B_{\nwarrow}(c_1)$, $A_{\uparrow}$, $A_{\leftarrow}$ and $A_{\rightarrow}$. Otherwise, the horizontal extent of $\mathcal{S}$ must be greater than $16\delta' r$, and we consider instead the sets $B_{\nearrow}(c_4)$, $B_{\nwarrow}(c_3)$, $A'_{\uparrow}$, $A'_{\leftarrow}$ and $A'_{\rightarrow}$. Here, $A'_{\leftarrow}$ and $A'_{\rightarrow}$ are rectangles of height $\delta' r$ and width $(1-4\delta')r$ placed to the left and right of cells c_3 and c_4 respectively, and $A'_{\uparrow}$ is any rectangle of height $(1-20\delta')r$ and width $16\delta' r$ glued on top of c_1 and strictly between the left side of c_3 and the right side of c_4. In both cases, we deduce that the topological component of $\bigcup(N(\mathcal{S}) \setminus \mathcal{S})$ that touches the upper piece of γ has area at least $(1+\delta'/6)\theta r^2/2$. Since some cells in this component touch one side of $[0,1]^2$, Lemma 5 implies that $N(\mathcal{S}) \setminus \mathcal{S}$ must contain some dense cell.

Case 3. Let $\mathcal{S} \subseteq \mathcal{C}$ be a connected set of cells which is not small. Suppose first that $N(\mathcal{S})$ touches exactly two sides of $[0,1]^2$ which are adjacent (say the bottom and the left sides of $[0,1]^2$). If the horizontal extent of $\mathcal{S}$ is at most $4r$, then $N(\mathcal{S}) \setminus \mathcal{S}$ has some cell inside the $5r \times 5r$ square on the bottom left corner of $[0,1]^2$. But these cells are all dense by Lemma 5 and we are done. Hence we can assume that $\mathcal{S}$ has horizontal extent greater than $4r$. In the other cases that $N(\mathcal{S})$ touches two non-adjacent sides or three sides of $[0,1]^2$, we can assume without loss of generality that $N(\mathcal{S})$ touches the left and right sides of $[0,1]^2$ but not the top side. Therefore, in all the cases considered, $\mathcal{S}$ must contain some cells intersecting each of the five first vertical stripes of width r at the left side of $[0,1]^2$. Let c_1, c_2, c_3, c_4 and c_5 be the uppermost cells in $\mathcal{S}$ intersecting each of the five vertical stripes. These cells are not necessarily all different, but for each c of these, either $B_{\nwarrow}(c)$ or $B_{\nearrow}(c)$ is completely contained in the corresponding strip. Thus, the topological component of $\bigcup(N(\mathcal{S}) \setminus \mathcal{S})$ that touches the upper piece of γ has area at least $5(1-4\delta')^2\theta r^2/4 > (1+1/8)^2\theta r^2$, and by Lemma 5, $N(\mathcal{S})\setminus\mathcal{S}$ must contain some dense cell.

Case 4. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be connected sets of cells not adjacent to each other (i.e. $\mathcal{S}_2\cap N(\mathcal{S}_1)=\emptyset$) and such that both $N(\mathcal{S}_1)$ and $N(\mathcal{S}_2)$ touch all four sides of $[0,1]^2$. Note that by Lemma 5 all cells inside the $5r\times 5r$ square on the top left corner of $[0,1]^2$ are dense. Assume that none of these cells belongs to $N(\mathcal{S}_1)\setminus\mathcal{S}_1$ or $N(\mathcal{S}_2)\setminus\mathcal{S}_2$ (otherwise we are done). It could happen that these cells in the top left square are either all in $\mathcal{S}_1$ or all in $\mathcal{S}_2$. Assume they are not in $\mathcal{S}_1$. Then consider, as in Case 3, the uppermost cells c_1, c_2, c_3, c_4 and c_5 in $\mathcal{S}_1$ intersecting each of the five first vertical stripes of width r at the left side of $[0,1]^2$. The same argument shows that the topological component of $\bigcup(N(\mathcal{S}_1)\setminus\mathcal{S}_1)$ that touches the upper left piece of γ has area at least $(1+1/8)^2\theta r^2$, and Lemma 5 completes the proof. □

Finally, we need to show that bad components a.a.s. have some properties to be used in the construction of the Hamilton cycles. Given a component b of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$, let $J=J(b)\subseteq\{1,\dots,n\}$ be the set of indices of points in $\boldsymbol{X}$ contained in some cell of b. Moreover, for any r', consider the set $J'=J'(b,r')=N_{\mathscr{G}(\boldsymbol{X};r')}(J)\setminus J$ (i.e. the set of strict neighbours of J in a random geometric graph of radius r').

Lemma 7. *For a small enough constant $\delta>0$, and for any r and r' satisfying $\theta nr^2\sim\log n$ and $r\le r'\le(1+1/32)r$, the following is a.a.s. true. For each small component b of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$, there exists a connected set of dense cells $\mathcal{R}(b)\subseteq\mathcal{D}_0$ of size $0<|\mathcal{R}(b)|\le 10/\delta^2$ such that*

1. *for every $i\in J'(b,r')$, the cell containing X_i is adjacent to some cell in $\mathcal{R}(b)$, and*
2. *$\mathcal{R}(b)\cap\mathcal{R}(\widetilde{b})=\emptyset$ and $J'(b,r')\cap J'(\widetilde{b},r')=\emptyset$, for any other small component $\widetilde{b}$ of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$. different from b.*

Proof. Let b be a small component of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$, and let g be any 16×16 grid covering b. Let O denote the geometric centre of the grid g, and let $\mathcal{S}$ be the set of cells which have some point at distance between $3r/4$ and $3r/2$ from O. Take as $\mathcal{R}(b)$ the subset $\mathcal{R}=\mathcal{S}\cap\mathcal{D}$ formed by the dense cells in $\mathcal{S}$. This set will be shown to have all the desired properties. (Note that the size of $\mathcal{R}$ is $|\mathcal{R}|\le|\mathcal{S}|<10/\delta^2$.)

Consider a coarser tessellation of $[0,1]^2$ into larger squares of side $\lfloor 1/(16\delta')\rfloor\delta' r$ (each square containing exactly $\lfloor 1/(16\delta')\rfloor^2$ cells). We refer to each square both as a subset of $[0,1]^2$, and as the set of cells it contains. Let $\mathcal{Q}$ be the set of squares of the coarser tessellation that contain at least one point at distance exactly $5r/4$ from O. By construction, all squares in $\mathcal{Q}$ are contained inside $\mathcal{S}$. Moreover, we claim that all squares in $\mathcal{Q}$ contain some dense cell. To show this, suppose first that $N(b)$ does not touch any side of $[0,1]^2$. In fact, by choosing δ sufficiently small, we can guarantee that each square $q\in\mathcal{Q}$ has no intersection with $N(b)\setminus b$, and thus $q\cup(N(b)\setminus b)$ is a connected set of cells of area at least

$$\theta(1-34\delta')^2r^2+\lfloor 1/(16\delta')\rfloor^2\delta'^2r^2\ge(\theta+1/257)r^2.$$

Hence, assuming that statement 1 in Lemma 5 holds, $q\cup(N(b)\setminus b)$ must contain some dense cell, which must be in q since $N(b)\setminus b$ does not contain any. The cases that $N(b)$ touches one or two sides of $[0,1]^2$ are dealt with analogously by using statements 2 and 3 in Lemma 5.

Since the union of squares in $\mathcal{Q}$ is topologically connected, and each pair of cells lying in topologically adjacent squares of $\mathcal{Q}$ are also adjacent in $\mathscr{G}_{\mathcal{C}}$, the dense cells in squares of $\mathcal{Q}$ induce a connected set of cells. Moreover, for any other cell c in $\mathcal{S}$ there is some square $q\in\mathcal{Q}$ such that c is adjacent to all cells in q. Hence, $N(\mathcal{R})\supseteq\mathcal{S}$, and also $\mathcal{R}$ induces a connected set of cells. Since $\mathcal{R}$ cannot be embedded in a 16×16 grid of cells, $\mathcal{R}$ must be contained in $\mathcal{D}_0$.

Now consider any vertex $i \in J' = J'(b, r')$. If $d(X_i, O) \le 3r/8$, then the cell c containing X_i must be in $N(b) \setminus b$. Therefore, since b is a component of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$, c must be sparse but adjacent to some dense cell $d \in \mathcal{D}_0$. By construction, any point in d must be at distance between $(1 - 34\delta')r$ and $(11/8 + 2\delta')r$ from O, so $d \in \mathcal{R}$. Otherwise, suppose that $d(X_i, O) > 3r/8$. We also have $d(X_i, O) \le (1 + 1/32 + 16\delta')r$, since $i \in J'$. Then the cell c containing X_i must be adjacent to all cells in some square $q \in \mathcal{Q}$, and in particular to some dense cell in $\mathcal{R}$.

To verify the other requirements, define $\mathcal{Q}'$ to be the set of squares of the coarser tessellation with some point at distance exactly $7r/4$ from O. The same argument we used for $\mathcal{Q}$ shows that all squares in $\mathcal{Q}'$ contain some dense cell. Let $\mathcal{R}'$ be the set of dense cells in squares of $\mathcal{Q}'$. Then it is immediate to verify that any point in a cell c of some other small component $\widetilde{b} \neq b$ of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$ must be at distance at least $41r/16$ from O since otherwise c would be adjacent to some cell in b, $\mathcal{R}$ or $\mathcal{R}'$. All remaining statements follow easily from that. $\square$

3 Packing linear forests in bipartite graphs

A *factorisation* of a graph is the set of subgraphs induced by a partition of the edge set. A *hamiltonian decomposition* of a graph is a factorisation in which at most one subgraph is a perfect matching, and all the remaining ones are Hamilton cycles. We call a matching that contains an edge of each of the Hamilton cycles in the decomposition a *transversal* of the decomposition. (Note that the transversal does not contain an edge of the perfect matching.) The construction of a hamiltonian decomposition in the following lemma is well known (since 1892). It is attributed to Walecki by Lucas [5]. We will use features of the construction in the proof of Lemma 2, and we use the transversal in Section 4.

Lemma 8. *Every complete graph has a hamiltonian decomposition with a transversal.*

Note that the number of Hamilton cycles in such a decomposition of K_{k+1} will be $\lfloor k/2 \rfloor$, and thus for k odd the transversal is not quite a perfect matching.

Proof. First, for k even, consider the complete graph K_{k+1} on the vertices $\{1, 2, \dots, k, *\}$. We shall first colour the edges of K_{k+1}. Expressions referring to vertex labels other than $*$ are interpreted mod k and expressions referring to colour labels are mod $k/2$. (In this paper, mod denotes taking the remainder on division.)

For each pair of vertices u and v in $\{1, 2, \dots, k\}$, assign the colour

$$\lceil (u + v)/2 \rceil \tag{3}$$

(mod $k/2$ of course) to the edge uv. Also, assign colour i to the edges from $*$ to both vertices i and $i + k/2$. See Figure 2.

It is easy to check that, for each $i \in \{1, \dots, k/2\}$, the edges receiving colour i form a $(k + 1)$-cycle $(v_0, \dots, v_k)$ where

$$v_0 = *, \qquad v_1 = i, \qquad v_{t+1} = v_t + (-1)^t t, \quad \forall t \in \{1, \dots, k - 1\},$$

or equivalently

$$v_0 = *, \qquad v_t = i - (-1)^t \lfloor t/2 \rfloor, \quad \forall t \in \{1, \dots, k\}.$$

Thus, the colouring induces a factorisation of K_{k+1} into $k/2$ Hamilton cycles of colours $1, \dots, k/2$, giving the required hamiltonian decomposition.

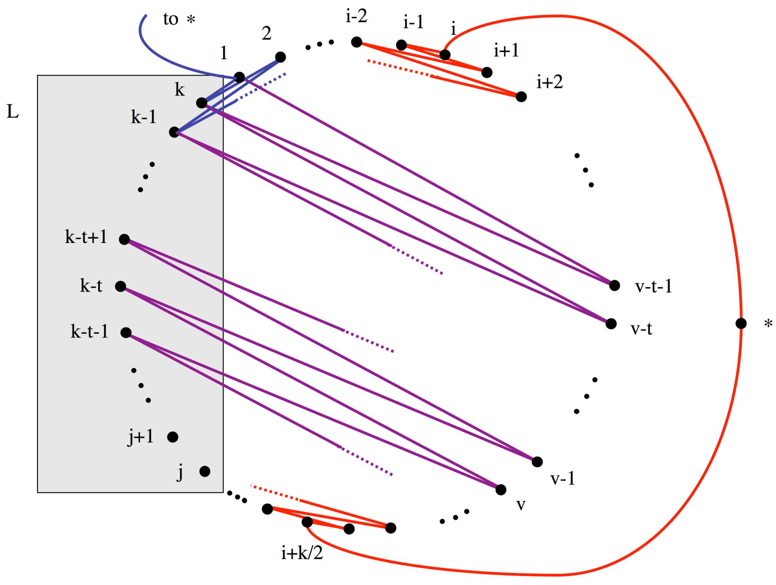

Figure 2: *Parts of a hamiltonian decomposition. The colour i is shown as red, colour 1 is blue, and colour $v+k-t$ is purple.*

When k is congruent to 2 mod 4, the set of edges $\{2i, 2i+1\}$ ($i = 0, \ldots, k/2-1$) is a transversal. When k is divisible by 4, one transversal uses the edges $\{2i, 2i+1\}$ ($i = 0, \ldots, k/4-1$), the edge from $*$ to $k/2$, and the edges $\{k/2+2i-1, k/2+2i\}$ ($i = 1, \ldots, k/4-1$).

For odd k, a perfect matching needs to be included. There is a similar colouring scheme, using the colours $1, \ldots, (k+1)/2$, where colours are taken mod $(k+1)/2$. In this case, colour $\lceil (u+v \mod k)/2 \rceil$ is on the edge uv (note we assume by convention that $u+v \mod k \in \{0, \ldots, k-1\}$), each colour i ($i \in \{1, \ldots, (k-1)/2\}$) is on the edge from i to $*$, and each colour $(k+1)/2 - i$ ($i \in \{0, \ldots, (k-1)/2\}$) is on the edge from $k-i$ to $*$. The edges of colour $(k+1)/2$ form a perfect matching, and each of the other colours gives a Hamilton cycle. Finally, the transversal for odd k is easily found, similar to the even k case, using more or less every second edge of the form $\{i, i+1\}$. □

Proof of Lemma 2 Since the statement is trivial for $j = 1$, we may assume that $j \geq 2$. Firstly, we deal with the case of even k. We will assign colours in $\{1, \ldots, k/2\}$ to a subset of the edges of G such that the colour classes determine the decomposition into linear forests. With a slight abuse of standard notation we will call this an edge colouring of G. For simplicity of exposition, we assume that the degree lower bounds for H are all met precisely. It will be evident from the proof that if any lower bounds on the degrees of vertices in J, that are specified in the lemma statement, are exceeded, it can only help by giving more choices in various steps. Hence, we may assume that the apex vertex has degree exactly $\ell + 1$, and all others have degree exactly ℓ. Label the apex vertex in J with '$*$', and label the other vertices $1, \ldots, j-1$.

First colour the edges of K_{k+1} as in the proof of Lemma 8, using the same labels on vertices as in that proof. If we delete the vertices in the set $L = \{j, j+1, \ldots, k\}$ of K_{k+1}, the remaining vertices of K_{k+1} are in J and the edges between them in G can inherit the colours from K_{k+1}. The set L should be regarded as a set of vertices sitting outside G. For $v \in J$, edges to other vertices in J are coloured, but the colour on any edge to L is 'missing' from that vertex in G. When we speak of missing colours, we count them with multiplicities: if a

v has no edge of colour i in G (i.e. two edges of colour i join v to L) then colour i is missing at v with multiplicity 2. Of course, each vertex is incident with two edges of each colour in K_{k+1}. The missing colours need to be assigned to edges of H, which go between J and B, since each vertex of J must have degree 2 in each of the final linear forests.

We will use a greedy colouring procedure to assign the missing colours to the edges of H, and thereby complete the desired edge colouring of G. The requirement is simply that each colour class must induce a linear forest. (By taking care of the missing colours we are ensuring that all vertices of J have two incident edges of each colour.) The procedure treats the vertices in the order $1, 2, \ldots, j-1$, and finally the apex vertex, $*$. Note that, so far, only edges within J are coloured, and each colour induces a set of paths. This is because the edges of any given colour induce a proper subgraph of the original Hamilton cycle of that colour. The procedure makes the assignments of new edges one by one, so it simply has to avoid all monochromatic cycles at each step, and terminate with each vertex in J having precisely two incident edges of each colour in G, just like they do in K_{k+1}, and each vertex in B $(= V(H) \setminus J)$ having at most two incident edges of any given colour.

The colouring procedure is defined inductively and requires several observations along the way. We will first fix a vertex $v \in J \setminus \{*\}$ (i.e. $1 \le v \le j-1$) and specify how the procedure treats the edges of H incident with v. We may assume inductively that there are no missing colours at vertices $1, \ldots, v-1$, i.e. all these vertices are incident with two edges of each colour, and furthermore that at each step the set of edges with any given colour induces a subgraph consisting of disjoint paths. The colours missing at v are those on the edges from v to L. Let i be a colour that is missing at v. There are two cases. In the first case, i is on just one edge at v, from v to $k-t$ say, and either $t = 0$ and no edges of colour i are already present in H, or $t = k - j = \ell - 1$. In the second case, i is on two edges, from v to $k-t-1$ and $k-t$, for some $0 \le t \le \ell - 2$. (In both cases of course t can be computed from the colour formula (3).) Let us refer to this colour i as i_t.

The specification of how the colouring procedure treats v is as follows. The missing colours i_t are treated one by one in decreasing order of t. At any point, let G_i denote the subgraph of H induced by the edges coloured i. Each missing colour i of multiplicity δ, in its turn, is assigned to δ uncoloured edges of H incident with v, in any manner such that:

(i) G_i remains a linear forest,

(ii) G_i is not a connected graph unless all choices satisfying (i) cause G_i to be connected.

Actually, we only need to invoke rule (ii) when $t = \ell - 3$, but it does no harm to enforce it in each step.

For the apex vertex $*$, which is treated last, the rule is simpler. The procedure assigns the missing colours to uncoloured edges of H incident with $*$, greedily subject to rule (i), and the order of treatment of the colours is determined at the start as follows: colours that already appear on more edges are treated earlier.

To verify that the colouring procedure must terminate with each colour inducing a linear forest, we argue inductively for $v \in J$, $v \ne *$, in the order of treatment. The apex vertex is considered last. We show that the linear forest condition holds after each vertex is treated. The argument for the inductive step also applies to the initial step, where $v = 1$.

So, for a vertex $v \in J \setminus \{*\}$, consider the point at which the procedure is treating the colour i_t defined above. Any edges coloured i_t at earlier steps of the inductive procedure must have been edges from $\{v-t-1, v-t, \ldots, v-1\} \setminus L$ to L, and these must go to the vertices $k-t+1, \ldots, k$. See Figure 2, where the colour i_t is shown in purple. Each vertex in L is incident with at most two such edges of colour i_t, as is each vertex in $\{v-t-1, v-t, \dots, v-1\} \setminus L$. In particular, $v-t-1$ is incident with exactly one such edge, the one joining it to vertex k. Hence, the number of times that colour i_t is missing (counted with multiplicities) on the vertices already treated is at most $2t+1$. Consequently, $2t+1$ is an upper bound on the number of edges of colour i_t in H at the point in the procedure where vertex v is about to be treated. It is not necessary, but may help to note that the distinct colours missing at v are either $i_0, i_2, \dots, i_m$ (where $m = 2\lfloor(\ell-1)/2\rfloor$), if i_0 has multiplicity 2, or $i_0, i_1, i_3, \dots, i_m$ (where $m = 2\lfloor \ell/2 \rfloor - 1$) if i_0 has multiplicity 1.

Recall that v has degree ℓ in H. Let δ denote the multiplicity of i_t as a missing colour at v. When the procedure is about to treat colour i_t at v, the colours on edges from v to $j, j+1, \dots, k-t-\delta$ have already been assigned. This means that $k-t-\delta-j+1$ of the edges incident with v are already assigned colours. As $d_H(v) = \ell = k-j+1$, there are precisely $t+\delta$ edges of H incident with v that remain uncoloured. Recall that H has at most $2t+1$ edges already coloured i_t. Since H is bipartite, at most t vertices in B can have degree 2 in G_{i_t}. Hence, there must be at least δ (which is either 1 or 2, precisely the multiplicity of i_t) uncoloured edges from v to vertices of B having degree at most 1 in G_{i_t}. So it is possible to assign colour i_t to δ of these edges at this point without creating any vertices of degree greater than 2 in G_{i_t}. We need to show that this can always be done so as to satisfy condition (i), that is, without creating a cycle in G_{i_t}.

Before proceeding, we need to understand the ways that such a cycle can form. If $\delta = 2$, then two uncoloured edges joining v to $N_H(v)$ must be picked with the purpose of colouring them i_t, and a cycle is created if and only if the two end-vertices in $N_H(v)$ are the two ends of a path in G_{i_t}. If $\delta = 1$, a cycle is created if and only if the edge picked is the end-vertex of a path in G_{i_t}, of which v is the other end-vertex.

Let U denote the set of vertices in $N_H(v)$ that are joined by uncoloured edges to v. It is always possible to avoid the cycle in question if U contains more than δ vertices of degree less than 2 in G_{i_t}. For, if this is true in the $\delta = 2$ case, then, even if one end of a path of G_{i_t} is picked for the first edge to be coloured i_t, the second edge can avoid the other end of the same path. The case $\delta = 1$ is even easier. There is similarly no problem if U contains a vertex of degree 0 in G_{i_t}.

So, if the the edges of colour i_t do not form a linear forest after the procedure treats v, we may assume that U contains precisely δ vertices of degree 1 in G_{i_t}, and that the other t vertices of U (recall that $|U| = t+\delta$) have degree 2 in G_{i_t}. In particular, the number of edges of colour i_t in H is precisely $2t+\delta$. However, above we deduced that it is at most $2t+1$. Hence, $\delta = 1$, and there must be precisely t vertices in B of degree 2 in G_{i_t}. That is, i_t is the colour of only one edge from v to L. As explained above, this happens only if, on the one hand, $t = 0$ and there are no edges already of colour i_t in H, or, on the other hand, $t = \ell - 1$. The first case contradicts the fact that U contains precisely δ vertices of degree 1 in G_{i_t}. So the second case holds, and t must be equal to $\ell - 1$. This means that u_t is the first colour being treated by the process at v, and furthermore, the $t+1$ vertices in U comprise exactly $N_H(v)$.

It was observed above that the only other vertices sending edges to L of colour i_t are $\{v-t-1, v-t, \dots, v-1\} \setminus L$ and the maximum number of such edges is $2t+1$. The edges in H of colour i_t come from these vertices. Since this upper bound is achieved, the situation is tight: $v-t-1$ is incident with one such edge in H, and the t vertices in $\{v-t-1, v-t, \dots, v-1\}$ are incident with two each. Recalling that U has precisely one vertex of degree 1 in G_{i_t} and the rest have degree 2, we conclude that the graph $G_{i_t} \cap H$ has precisely two vertices of degree 1 and the rest of degree 2. Since by induction it contains no cycle, it is a path P, and one end-vertex of P is $v-t-1$. Considering the original colouring of J, the only other edges of colour i_t in G at this point form a path on the vertices $\{1,\dots,v-t-1\}$ that is vertex-disjoint from P except for the vertex $v-t-1$. Hence, G_{i_t} is a path and is hence connected.

Shortly we will also need a different observation. Note that if $t=0$, which means $\ell=1$, then the hypotheses imply that v has precisely one neighbour in B, which is distinct from the neighbours of $1,2,\dots,v-1$. So this cannot be the case. It follows that $t+1=\ell\geq 2$.

Since there is at least one edge already coloured i_t, we have $v\geq 2$. Hence, the vertex $v-1$ was already treated, and has two edges of this same colour i_t in H joining it to B. Call them x and y. Shift the focus back to the time that the procedure, when treating $v-1$, coloured x and y. After they were coloured, as observed above, G_{i_t} became connected. By rule (ii), it was necessary that that no other choice of two uncoloured edges from $v-1$ to B could avoid creating a connected graph G_{i_t}. However, x and y are two consecutive edges in the path P defined above. Let w denote any vertex of $N(v-1)\cap B\setminus N(v)$, which must exist by the second hypothesis of the lemma. Note that w is adjacent to no edges of colour i_t. Also note that i_t is the first colour treated by the procedure when dealing with the vertex v_{t-1}, since this is the colour of the edge from v_{t-1} to the vertex $j=k-t$. Hence, at that point in the procedure, all edges of H incident with v_{t-1} are uncoloured, in particular the edge to w.

We consider two cases. Firstly, if x or y is incident with an end-vertex u of P, it must be that u lies in B. There are no other edges of colour i_t incident with u. So, instead of colouring x and y using the colour i_t, the edges to u and w could have been used instead, and these would form a separate path in G_{i_t}, making it disconnected, which is a contradiction by rule (ii). If, on the other hand, x and y are elsewhere in P, then the graph P' induced by the edges of P other than x and y is disconnected. In this case, the procedure could have placed the colour i_t on the edge x, and on the edge from v_{t-1} to w, again a contradiction since G_{i_t} becomes disconnected.

It follows that our assumptions about U above (that it contains precisely δ vertices of degree 1, etc.) are false. This concludes what was required to show that after the procedure treats $v\in J\setminus\{*\}$, the edges of each colour induce a linear forest.

Finally, we turn to the apex vertex, $*$. Note that the multiset of colours missing at $*$ is precisely $\{j,j+1,\dots,k\}$: each such colour i lies on an edge from $*$ to the vertex i.

We first consider $j>k/2$. In this case the missing colours at $*$ all have multiplicity 1, so it is a simpler situation than for smaller j. If we list the colours in the following order: $i_1=j$, $i_2=k$, $i_3=j+1$, $i_4=k-1$, $i_5=j+2$ and so on, it is easily seen that the number of edges already coloured i_t is precisely $2\ell+1-2t$ $(1\leq t\leq \ell)$. So the colouring procedure treats them in the order $i_1,i_2,\dots,i_\ell$. Since $*$ has degree $\ell+1$ in H, the number of uncoloured incident edges it has when treating colour i_t is $\ell+2-t$. At this point, there are only enough edges already coloured i_t to create at most $\ell-t$ vertices of degree 2 in B. Hence, there are still at least another two vertices in B joined to $*$ by uncoloured edges.The only way to create a cycle in colouring one of these edges i_t is to use the edge that joins to the other end of the unique path in G_{i_t} that presently begins with $*$. Hence, there is yet another edge available to safely colour i_t so as to maintain the linear forest condition.

It only remains to show that a similar statement holds for the case $j\leq k/2$. Here, the colours $1,\dots,j-1$ are missing with multiplicity 1 at $*$, and the colours $j,\dots,k/2$ are missing with multiplicity 2. At this point in the colouring procedure, the numbers of edges of colours $1,j-1,2,j-3,\dots$ are $2j-3,2j-5,2j-7,\dots$ respectively. So the procedure will treat these colours in that order. On the other hand, for each $j\leq i\leq k/2$, the number of edges of colour i is precisely $2j-2$, and these were placed in pairs, two at a time from each of the vertices $1,\dots,j-1$. So the procedure treats these colours first. When treating colour $i\in\{j,\dots,k/2\}$, the number of uncoloured edges from $*$ to B will still remain at least $\ell + 1 - (k - 2j) = j + 2$ before each colouring step. With only $2j - 2$ edges already of colour i in H, it is easy to assign two more edges of H incident with $*$ the colour i without creating a cycle or vertex of degree bigger than 2 in G_i: either G_i is a single path with j vertices in B, in which case there are two edges to vertices of degree 0 in G_i, or G_i has at least two components, in which case the procedure can avoid the at most $j - 1$ vertices of degree 2 in G_i and join instead to two of degree 1 or 0, and not in the same component of G_i. For the remaining colours in $\{1, \ldots, j - 1\}$, set $i_1 = 1$, $i_2 = j$, $i_3 = 2$ etc in the order given above, so that the number of edges of colour i_t is $2j + 1 - 2t$. So, the argument as in the case $j > k/2$ (with ℓ replaced by j in the appropriate places) shows that the colours can be assigned as required.

Finally, we consider the case of odd k, which will be reduced to the even case. Note that the colouring in Figure 2 is rotationally symmetric, and thus in the even k case, we can apply the same colouring procedure if the labels of the vertices of the auxiliary K_{k+1} are shifted by some quantity (i.e. we can use the labels in $\{a+1, \ldots, a+j-1, *\}$ for the vertices in J and delete from K_{k+1} the vertices in $L = \{a + j, a + j + 1, \ldots, a + k\}$). So, for the case of odd k, label the apex vertex in J with $*$, and label the other vertices $\lceil \ell/2 \rceil, \ldots, \lceil \ell/2 \rceil + j - 2$. Then, colour the edges of an auxiliary K_{k+1} as in the proof of Lemma 8, using the same labels on the vertices as in the proof. We refer to the colour $(k + 1)/2$ as the *match* color. Delete from K_{k+1} the vertices in the set $L = \{\lceil \ell/2 \rceil - i \mod k \ : \ 1 \le i \le \ell\}$, and identify the set of remaining vertices with J. As in the case of even k, we must assign all the missing colours to the edges of H with the same requirements as before on the non-match colours, but also the match colour must induce a matching in H. Note that this last condition is trivially satisfied since, by our specific choice of J and L, at most one edge coloured with the matching colour crosses between J and L.

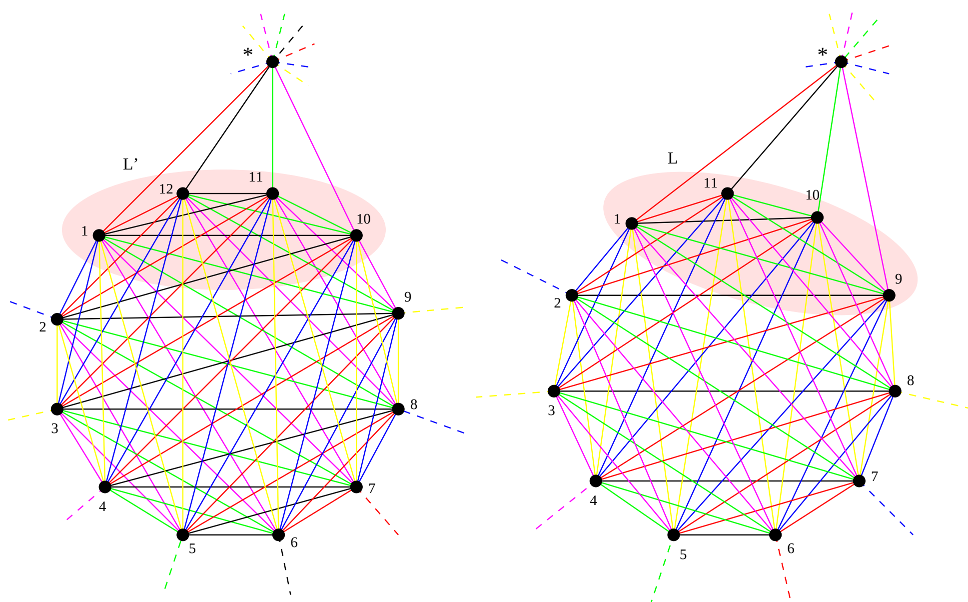

Figure 3: *Reduction of the case $k = 11$ to the case $k = 12$.*

To find the required colour assignment to the edges of H, consider the bipartite graph H' with parts J' and B resulting from adding a new vertex labelled $\lceil \ell/2 \rceil + j - 1$ to J, with ℓ edges from the new vertex to arbitrary vertices of B. We shall use the previous greedy procedure to colour the edges of H' noting that $|J'| + \ell - 1 = k + 1$ is even. Take a copy of K_{k+2} (disjoint from K_{k+1}) with labels $1, \ldots, k+1, *$ on the vertices, and colour the edges as in the proof of Lemma 8. Partition the vertex set of K_{k+2} into $J' = \{\lceil \ell/2 \rceil, \dots, \lceil \ell/2 \rceil + j - 1, *\}$ and $L' = \{\lceil \ell/2 \rceil - i \mod k+1 \ : \ 1 \le i \le \ell\}$, and delete the vertices in L'. Colour the edges of H' using the previously described colouring procedure (with all vertex labels shifted by the appropriate constant, so that the vertices in L' have the correct labels). Even though the edge colours used in K_{k+1} and K_{k+2} are different, it is easy to check that, for each $i \in \{\lceil \ell/2 \rceil, \dots, \lceil \ell/2 \rceil + j - 2, *\}$, the vertex of J that is labelled i misses exactly the same set of colours as the vertex of J' labelled i (see Figure 3 for a visual illustration). Therefore we may simply obtain the edge colouring of H' from the algorithm used for even k applied to H', and then restrict this colouring to H to obtain a colouring that satisfies the desired properties. □

4 Building Hamilton cycles and a perfect matching

In this section, we use the results in the earlier lemmas to prove Theorem 1. We first give a complete proof for k even and then provide the extra pieces of argument required for k odd.

(i) **Proof for k even.**

Let $\epsilon > 0$ be arbitrarily small. Recall the definitions of m and r_k in the beginning of Section 2. In view of Proposition 3, we can choose a large enough constant $\lambda > 0$ such that, by setting

$$r_l = \sqrt{\frac{\log n + m \log\log n - \lambda}{\theta n}} \qquad \text{and} \qquad r_u = \sqrt{\frac{\log n + m \log\log n + \lambda}{\theta n}},$$

we can guarantee that $\mathbf{Pr}(\mathscr{G}(\boldsymbol{X}; r_l)\ k\text{-connected}) < \epsilon/2$ and $\mathbf{Pr}(\mathscr{G}(\boldsymbol{X}; r_u)\ k\text{-connected}) > 1 - \epsilon/2$. In other words, looking at the evolution of $\mathscr{G}(\boldsymbol{X}; r)$ for $0 \le r < \infty$, the probability that it becomes k-connected at some point r_k between r_l and r_u is greater than $> 1 - \epsilon$. Let us condition upon this event. By the results in Section 2, we may assume that the properties described in Lemmas 6 and 7 hold for $r = r_l$ and some δ, and also that the property in Lemma 4 holds for $r = r_k$ and $\eta = 32\delta$. So we may assume $\boldsymbol{X}$ to be an arbitrary fixed set of n points in $[0,1]^2$ in general position and satisfying these properties. The proof is completed by giving a deterministic construction of $k/2$ edge-disjoint Hamilton cycles for the geometric graph $\mathscr{G}(\boldsymbol{X}; r_k)$. Most edges will be of length at most r_l but we shall use a few of length between r_l and r_k. (The last edges creating k-connectivity arrive during this period, and they are of course necessary to construct $k/2$ edge-disjoint Hamilton cycles.) We define the edges of each Hamilton cycle by colouring some of the edges of $\mathscr{G}(\boldsymbol{X}; r_k)$, using colours $1, \dots, k/2$, such that each of these colour classes induces a Hamilton cycle.

We take $r = r_l$ (except at special points in the argument) and define $\mathscr{G}_{\mathcal{C}}$, $\mathcal{D}$, $\mathcal{B}$ and so on accordingly (see Section 2). Let T be a spanning tree of the largest component $\mathcal{D}_0$ of $\mathscr{G}_{\mathcal{C}}[\mathcal{D}]$. Next, double each edge of T to get an Eulerian multigraph F. The vertex degrees in T are bounded above by Δ, so those in F are bounded above by 2Δ. Next, pick an Eulerian circuit C of F.

Henceforth, we have no need to consider points in $[0,1]^2$ that are not members of $\boldsymbol{X}$. So, points in $\boldsymbol{X}$ contained in some cell c will simply be referred to as points in c, and they will be often identified with their corresponding vertices in $\mathscr{G}(\boldsymbol{X}; r_l)$ or $\mathscr{G}(\boldsymbol{X}; r_k)$. Also, the term *dense cell* will refer only to cells in $\mathcal{D}_0$, thus excluding these dense cells contained in bad components. For descriptive purposes, we split the rest of the argument into two parts, first treating the case that there are no bad cells, i.e. $\mathcal{B}$ is empty. For this we only need the edges of $\mathscr{G}(\boldsymbol{X}; r_l)$. Then we will show how the construction is easily modified to handle the bad components, using some edges of $\mathscr{G}(\boldsymbol{X}; r_k)$.

Part 1. $\mathcal{B}$ is empty.

In this case, the rest of the proof involves two steps, which will be used in different forms during the later arguments.

Step 1. Turning the circuits into cycles

The subgraph of $\mathscr{G}(\boldsymbol{X}; r_l)$ induced by the points contained in any dense cell is complete and has many more than k vertices. Lemma 8 provides $k/2$ edge-disjoint Hamilton cycles in this subgraph. In fact, it provides more; we just choose a subset of the Hamilton cycles that are given by that lemma. The separate cycles in all the dense cells will be 'broken' and rejoined together using C as a template. In the following discussion we assume C is oriented, so we may speak of incoming and outgoing edges of C with respect to a cell.

For any dense cell c, the deletion of c from C breaks C up into a number of paths P_i. For colour 1, do the following. Associate each path P_i with an edge z_i that joins two points in c and has already been coloured 1, using a different edge z_i for each path P_i. Uncolour the edges z_i, and associate the outgoing and incoming edges of the path P_i (with respect to the cell c) each with an endpoint of z_i. After doing the same for all dense cells, every edge cd of C, where c and d are cells, has now been associated with two points, one in c and one in d. Colour the edge joining these two points using colour 1. Doing this for all edges of C clearly joins up all the edges coloured 1 into one big cycle using all points in the dense cells.

Now do the same with colours $2, \ldots, k/2$, one after another, but each time being careful to use edges z_i in each cell that are not adjacent to such edges used with any of the previous colours. This is easily done because using an edge for one colour eliminates at most four edges of another colour (as the graph induced by edges of a given colour has maximum degree at most 2). So the process can be carried out if M is greater than $2k\Delta$.

Step 2. Extending the cycles into the sparse cells

There are now $k/2$ edge-disjoint coloured cycles, one of each colour, and each cycle uses precisely all the points in dense cells. Note that within each dense cell, there are still an arbitrarily large number (depending on M) of *spare* edges of each colour, left over from the original application of Lemma 8. To prepare for extending the cycles into the sparse cells, we will break the cycles at these spare edges.

Let c be any sparse cell. By the definition of $\mathcal{B}$ and our assumption that $\mathcal{B}$ has no cells, there is a dense cell, say c', adjacent to c in $\mathscr{G}_{\mathcal{C}}$. If c contains at most $2k$ points, for each vertex v of the geometric graph inside c do the following. Choose a spare edge z inside c' of colour 1, uncolour the spare edge z, and colour the two edges from the endpoints of z to v with the colour 1. Any edges of different colours adjacent to z should be deemed not spare after use. Then repeat for each of the other colours. After this, the edges of any given colour form a cycle containing all points in dense cells and in c.

On the other hand, if c contains more than $2k$ points, the above process could potentially require too many spare edges, so we must do something else. By Lemma 8, we can specify $k/2$ edge-disjoint Hamilton cycles around the points in c, one of each of the colours. One can then greedily choose an independent set of edges, one of each colour. (This is easily seen by noting that choosing an edge knocks out at most four adjacent edges with any particular colour. Alternatively, by a more careful argument which we give later, it can be shown that the same holds as long as c contains more than k points.) These edges can be matched up with $k/2$ spare edges that have both endpoints in c', and then each of the coloured cycles is easily extended by uncolouring each matched pair of edges and appropriately colouring the edges joining their endpoints. Again for this case, the edges of any given colour form a cycle containing all the points in dense cells and in c.

This process can be repeated for each sparse cell. Since each dense cell has at most Δ neighbours in $\mathscr{G}_{\mathcal{C}}$, the total number of spare edges required of any one colour in any dense cell can be crudely bounded above by $2k\Delta$, which is the same as the upper bound on the number of points already used up. Thus, for M sufficiently large ($4k\Delta$ should do), there will be a sufficient number of spare edges to finish with a cycle of each colour through all points in $\mathscr{G}(\boldsymbol{X}; r_l)$, using only the edges of $\mathscr{G}(\boldsymbol{X}; r_l)$. This finishes Step 2 and the proof in the case that $\mathcal{B}$ is empty.

We now turn our attention to the (much more difficult) case that $\mathcal{B}$ is nonempty, for which we use an appropriate modification of the above argument.

Part 2. The general case: $\mathcal{B}$ can be nonempty.

For this, we will need to use some edges of $\mathscr{G}(\boldsymbol{X}; r_k)$ that are not present in $\mathscr{G}(\boldsymbol{X}; r_l)$, but the definition of all structures (such as bad components) remains as determined by the graph $\mathscr{G}(\boldsymbol{X}; r_l)$. Recall the Eulerian circuit C chosen at the start of the proof. This circuit gives a directed tour of all dense cells in the graph $\mathscr{G}_{\mathcal{C}}$. We will first extend it to a circuit that includes routes through each bad component, and later perform modified versions of Steps 1 and 2 described above.

Pick one such bad component b, which must be small by Lemma 6, and let $\mathcal{R} = \mathcal{R}(b)$ be a set of cells as in Lemma 7. Recall that $0 < |\mathcal{R}| < 10/\delta^2$. To take care of b, we will work entirely in $\mathcal{R}$ and the bad component b. Let J denote the set of points in cells in b and set $j = |J|$ (assume that $j > 0$, since otherwise b has no role in our argument). The subgraph of $\mathscr{G}(\boldsymbol{X}; r_l)$ induced by J is a copy of K_j, since b is small and can be embedded in a 16×16 grid of cells (and assuming that $32\delta < 1$). Now consider the graph $\mathscr{G}(\boldsymbol{X}; r_k)$, which by definition is k-connected, and let $J' = N_{\mathscr{G}(\boldsymbol{X}; r_k)}(J) \setminus J$. Let H denote the induced bipartite subgraph of $\mathscr{G}(\boldsymbol{X}; r_k)$ with parts J and J', and let $G \subseteq \mathscr{G}(\boldsymbol{X}; r_k)$ be the union of H with the clique on vertex set J. Note for later reference that the set J' can possibly contain vertices in dense cells: although no cell in b is adjacent to a dense cell, points in it can be adjacent to points in a dense cell.

Claim 1. G contains $k/2$ pairwise edge-disjoint linear forests $F_1, \dots, F_{k/2}$, such that in each forest
(a) all vertices in J have degree 2, and
(b) at most $2k$ vertices in J' are contained in any path.

To prove the claim, we consider two cases, the second being much more difficult since it requires Lemma 2. Later in Section 5 we provide an alternative proof of the claim which does not use Lemma 2.

Case 1: $j > k$.

Since $\mathscr{G}(\boldsymbol{X}; r_k)$ is k-connected, no vertex cut of G of size less than k can separate J from J'. Moreover, both J and J' have cardinality at least k. So (a version of) Menger's theorem implies that there is a set of k pairwise disjoint paths joining J to J'. Hence, there is a matching, T, of cardinality k, with each edge of the matching joining a point in J to a point in J'.

Consider first an arbitrary complete graph K_j, of which Lemma 8 can be used to obtain a full hamiltonian decomposition, together with a transversal containing one edge from each of the Hamilton cycles. Now choose $k/2$ of the Hamilton cycles in the decomposition, and let T' be the matching consisting of the edges of the transversal that lie in the chosen cycles.

Next, we can identify the set of vertices of K_j with the set J, such that the vertices incident with edges in T' are identified with the vertices of J that are matched by T. From each of the Hamilton cycles, delete the edge, say x, in that cycle that lies in T', and add the two edges of T adjacent to x. This gives a path P in G which starts and finishes in vertices in J'. Since T' is a matching, the end vertices of all paths comprise a set of k distinct vertices. Hence, these $k/2$ paths suffice for $F_1, \ldots, F_{k/2}$.

Case 2: $j \leq k$.

Since $\mathscr{G}(\boldsymbol{X}; r_k)$ is k-connected, each vertex in J has degree at least k. The case $j = 1$ is trivial, so we restrict our attention to $2 \leq j \leq k$. By Lemma 4, we may assume that one vertex in J, say X, has degree at least $k+1$. Hence, we have $d_H(X) \geq \ell + 1$ where $\ell = k - j + 1$, and $d_H(Y) \geq \ell$ for all other vertices Y in J. Moreover, if any two vertices X and Y in J had at most ℓ neighbours in J' jointly, then those neighbours, together with the remaining $j - 2$ members of J, would form a $(k-1)$-vertex cut, a contradiction. So the second condition in Lemma 2 is satisfied. Thus, Lemma 2 ensures the existence of a colouring of (some of) the edges of G, such that each colour induces a linear forest in which each vertex in J has degree 2 and whose ends are in J'. Since all edges in the paths of these forests are incident with J, at most $2j < 2k$ vertices in J' are used for each forest. So these forests can serve as the requisite linear forests $F_1, \ldots, F_{k/2}$. This completes the proof of Claim 1.

Let J'' be the set of vertices in J' contained in some path of $F_1, \ldots, F_{k/2}$. By the claim above, we have $|J''| < k^2$. Moreover, by setting $r = r_l$ and $r' = r_k$ in Lemma 7, we deduce that each vertex in J'' is contained in a cell c that is either dense or adjacent to some dense cell in $\mathcal{R}$. (Note that c may be either sparse or dense.) We extend each forest F_i to a spanning forest F'_i of $J \cup J''$ by adding those vertices in J'' not used by any path of F_i as separate paths of length 0. The total number of paths is at most $|J''|k/2 < k^3/2$, taking into account multiplicity since each 0-length path may belong to more than one forest F'_i.

The next step is to associate each of the paths in F'_i with the colour i, and create circuits $C_1, \ldots C_{k/2}$ such that circuit C_i contains C, together with an extra cycle $C(P)$ for each path P in the forest F'_i. To construct $C(P)$, take two cells c and d in $\mathcal{R}$ (possibly $c = d$), one for each end-vertex of P, either containing the end-vertex or adjacent to a cell containing it. Then the cycle $C(P)$ consists of a special new edge (possibly a loop) joining cells c and d (we say this edge *represents* P), together with a path of cells within $\mathcal{R}$ joining those same cells c and d. Note that each cycle has length bounded above by $10/\delta^2$ (see Lemma 7).

This construction was all with respect to a particular bad component b. Now repeat the construction for all the other bad components, in each case extending the circuits $C_1, \ldots C_{k/2}$ in the same way as for b. By Lemma 7, two paths P and P' related to different bad components have no vertex in common, and also the corresponding extra cycles $C(P)$ and $C(P')$ use disjoint sets of dense cells. Hence, the number of these new cycles passing through any particular dense cell is at most $k^3/2$, so the maximum degree of dense cells in each C_i is at most $2\Delta + k^3$.

We next perform a version of Step 1 described in Part 1. First, let us call all the vertices lying in forests F_i with respect to bad components the *forest* vertices.

Step 1'. Turning the circuits C_i into edge-disjoint coloured cycles

The first part of this is done just as for Step 1 when $\mathcal{B}$ empty, except for two aspects. To start with, all forest vertices within dense cells are set aside and not used in construction of the coloured cycles within the dense cells. Secondly, where an edge of the circuit between cells

c and c' is one of those representing a path P in a forest, instead of a simple edge between vertices u in c and v in c', the cycle uses the path P represented by that edge, together with the edges joining P to u and v. In this way, we obtain from C_i a cycle of colour i that visits precisely all vertices that are either forest vertices (i.e. in $J(b)$ or $J''(b)$ for any small component b), or in a dense cells, but *not* both. All other points in $\boldsymbol{X}$ will be called *outsiders.* They are not yet visited by any of $C_1, \ldots, C_{k/2}$, either because they are neither forest vertices nor in a dense cell, or are in both.

We next perform a version of Step 2, as follows.

Step 2'. Extending the cycles to the outsiders

This consists of extending each coloured cycle as done in Step 2, but this time extending to them only through the outsiders. The other significant difference between this and Step 2 is that the maximum degree of dense cells in each C_i is now bounded above by $2\Delta + k^3$ rather than 2Δ, and there are some outsiders in each dense cell, so there are fewer spare edges to work with, but the change is only a constant. So we need to adjust the lower bound on M accordingly. This completes the proof of part (i) of the theorem.

(ii) **Proof for k odd.**

We now consider odd $k \geq 1$. Recall that the total number of vertices in the geometric graph must be even. The same framework of argument is used as for k even. For k odd, colours $1, \ldots (k-1)/2$ will be used for the Hamilton cycles, and colour $(k+1)/2$ will be used for the matching. We find it convenient to refer to $(k+1)/2$ as the *match* colour. The edges that we colour using the match colour will form a spanning subgraph D of the geometric graph, whose edges form a cycle on some (possibly all) vertices and a matching of some of the other vertices; in the very last stage of the argument we will adjust this to form a perfect matching of the whole graph.

Define the multigraph F as for k even, and construct C in the same way. We next need to perform a version of Step 1. In this case, instead of creating $k/2$ edge-disjoint cycles passing through all the points in dense cells, we construct $(k+1)/2$ of them using the same construction as for k even.

In the case that there are no bad cells, the argument as in Part 1 above shows that all the cycles can be extended as in Step 2 to edge-disjoint Hamilton cycles of the graph. Then, since the graph has an even number of vertices, every second edge of the matching colour can be omitted to provide the desired colouring.

So consider the case that $\mathcal{B}$ is possibly nonempty and follow the argument in Part 2 for k even, up to the point where Claim 1 is made. This claim is replaced by the following.

Claim 2. G contains $(k+1)/2$ pairwise edge-disjoint linear forests $F_1, \ldots, F_{(k+1)/2}$, such that in each of the first $(k-1)/2$ forests

(a) all vertices in J have degree 2, and

(b) at most $2k$ vertices in J' are contained in any path,

and the last forest, $F_{(k+1)/2}$, is a matching that saturates each vertex in J.

To prove this claim, we again distinguish two main cases but insert an extra one. We also give an alternative proof in Section 5 which does not require Lemma 2.

Case 1a: $j > k + 1$.

As in the case $j > k$ for the first claim, we may find the matching T of (odd) cardinality k, and also $(k+1)/2$ edge-disjoint Hamilton cycles in $G[J]$, with a matching T' containing one edge from each Hamilton cycle. By relabelling J, we can then align the end vertices of the T' with those of T in J, but only using one end vertex of the edge in the last cycle. A set of edges in this last cycle is easily deleted so that what remains, together with the incident edge of T, is a matching either entirely contained in J or with one edge (of T) leaving J. This proves the claim in this case.

Case 1b: $j = k+1$.

In this case, we may use the decomposition of a $K_{j+1} = K_{k+2}$ into $(k+1)/2$ Hamilton cycles, delete any one vertex, and use the broken Hamilton cycles for the paths, plus every second edge of any one of them for the matching.

Case 2: $j \le k$.

This is identical to the proof of Case 2 of the first claim, using the appropriate part of Lemma 2. This finishes the proof of Claim 2.

We now extend the forests to spanning forests F_i', exactly as for the even k case (this now includes the match colour). We are now prepared for the analogue of Step 1.

Step 1". Circuits into cycles and matching

For each non-match colour $1 \le i \le (k-1)/2$, the forests F_i' in the various bad components are treated the same way as in Step 1' in order to create circuits C_i. These are then turned into edge-disjoint cycles $\widetilde{C}_i$ passing through all vertices except for the outsiders, just as for k even. Simultaneously with this, by including an extra colour in the construction, we create a cycle in the match colour that passes through just the vertices in the dense cells apart from outsiders. For this colour, we must do something different with respect to the vertices not in dense cells. So let $i = (k+1)/2$; so far we have a cycle $\widetilde{C}_i$, through all non-forest points in dense cells, whose edges are of colour i. The edges of a given forest F_i related to a bad component are naturally coloured with i. For each single vertex component v in the forest F_i' (i.e. each vertex v unmatched by F_i), let d be a dense cell either equal or adjacent to the cell c containing v. We may select the end-vertices v_1 and v_2 in d of a spare edge of $\widetilde{C}_i$ (note this implies that the edges vv_1 and vv_2 are currently uncoloured). Denoting these two vertices v_1 and v_2, the pair $\{v_1, v_2\}$ is called the *gate* for v. Each such vertex v is treated in this way, its gate is defined and v is added to a set W (a set of vertices which are 'waiting'). Naturally, any spare edge incident with either vertex in a gate is deemed non-spare for all subsequent choices of gates. Note that at this stage, all vertices in the graph are incident with an edge of the match colour except for the outsiders and those in W.

Next, we perform the step of extending the cycles to the outsiders.

Step 2". Extending cycles and matching to outsiders

For each non-match colour i, the cycle of colour i is extended as in Step 2'. We next show that we can also include a near-perfect matching of the match colour, which saturates all but a bounded number of outsiders, such that the matching is edge-disjoint from all the coloured cycles. This extra matching is easily obtained using the methodology of Step 2': for the match colour i, if a sparse cell c contains at most $2k$ outsiders, we may leave all outsiders unmatched. If it contains more outsiders, simply include an extra cycle $\widehat{C}$ through all outsiders in c. This cycle should be chosen simultaneously with all the other cycles being chosen within cell c in this Step 2", using Lemma 8. Then, colour every second edge of $\widehat{C}$ with the match colour, leaving at most one outsider in this cell unmatched. Finally, in either case, for each remaining unmatched outsider v, choose a gate (v_1, v_2) exactly as described in Step 1". Note that a dense cell contains a bounded number of outsiders since these are all in the forest F_i.

After all this, every vertex is in each of the coloured cycles of colour $i \le (k-1)/2$, but we still need to create the perfect matching of colour $i = (k+1)/2$. So far, all vertices are either matched by the match colour i, or lie in W, or lie on the cycle $\widetilde{C}_i$. To fix this, in one fell swoop, we choose simultaneously for all vertices v in W, a vertex v', in the gate for v, such that all the vertices v' have odd distance apart as measured along the cycle $\widetilde{C}_i$ of colour i. (Why this is possible will be explained shortly.) Then all such edges of the form vv' are coloured i, and finally, every second edge along C_i between these vertices is coloured i in such a way as to create a matching. The edges of colour i clearly form a perfect matching of the graph.

The only thing left to explain is why the choice of all v' as specified, creating odd distances, is feasible. Since there are two adjacent vertices on the cycle in each gate that can potentially be used as v', we may pass along the cycle $\widetilde{C}_i$ making sure that the distances between chosen vertices are odd, until returning to the starting point. The very last distance must be odd because the number of gates equals the number of vertices in W. These are precisely the vertices outside the cycle that are not already matched by colour i. Since the number of vertices in the graph is even, the parity is correct for every distance to be odd. □

We note here how the proof would be affected if Conjecture 1 were established. Note that condition (ii) in Lemma 2 is implied by k-connectedness, as is condition (i') in the conjecture. It is clear from our proof that it would then follow that the 'small' bottlenecks do not exist when the graph is k-connected, and the probability the graph does not have the required set of cycles (or cycles plus a matching) is much smaller than $1/\log n$. We don't investigate this further at this point as the conjecture is unproved.

5 Alternative for the packing argument

In this section we show how the proof of Theorem 1 can be shortened so that it does not require Lemma 2. First we provide a probabilistic argument that shows that 'dangerous' sets of vertices with relatively few common neighbours are rare, thus discarding the most difficult situations that Lemma 2 had to deal with. The following result implies Lemma 4 and has a fairly similar proof, so we omit some of the details. Recall the definition of m in Section 2.

Lemma 9. *For any small enough constant $\eta > 0$ and any r such that $\theta n r^2 = \log n + m \log\log n + O(1)$, and fixed $j \geq 2$, the random geometric graph $\mathscr{G}(\boldsymbol{X}; r)$ a.a.s. satisfies the following property. Every set J of j vertices such that $\max_{u,v \in J}\{d(X_u, X_v)\} \leq \eta r$ has at least k common neighbours.*

Proof. Given fixed integers $j \geq 2$, $0 \leq q \leq k-1$ and $s \geq 0$, we say that vertices $v_1, \ldots, v_{j+q+s}$ form a *bad* configuration if the following conditions hold: points X_{v_1} and X_{v_2} are at Euclidean distance $\rho \leq \sqrt{2}\eta r$; X_{v_1} is closer to the boundary of $[0,1]^2$ than X_{v_2} is; $X_{v_3}, \ldots, X_{v_j}$ lie in the circle with centre at X_{v_1} and radius ρ; vertices $v_{j+1}, \ldots, v_{j+q}$ are common neighbours in $\mathscr{G}(\boldsymbol{X}; r)$ of all $v_1, \ldots, v_j$; and $v_{j+q+1}, \ldots, v_{j+q+s}$ are neighbours in $\mathscr{G}(\boldsymbol{X}; r)$ of some but not all of the vertices $v_1, \ldots, v_j$. (The requirement on X_{v_1} and X_{v_2} is expressed in terms of the Euclidean distance only for computational purposes, regardless of the ℓ_p norm we use in the construction of $\mathscr{G}(\boldsymbol{X}; r)$, and it implies in particular that $d(X_{v_1}, X_{v_2}) \leq \eta r$.) It is easy to verify that, if the a.a.s. property in the statement fails, then we must have some bad configuration of vertices. Hence, it suffices to prove that the expected number of such bad configurations is $o(1)$. The computation of this expectation is very similar to that in the proof of Lemma 4 and is thus omitted. □

Note that the proof of Theorem 1 in Section 4 uses Lemma 2 only to deduce Claims 1 and 2. Therefore, we just need to re-prove these two claims using Lemma 9 rather than Lemma 2.

Alternative proof of Claim 1. First observe that we can find a set J'' of k points in J' which are common neighbours of all points in J with respect to $\mathscr{G}(\boldsymbol{X}; r_k)$. This follows from Lemma 9 for $j \geq 2$, and is trivially true for $j = 1$ since $\mathscr{G}(\boldsymbol{X}; r_k)$ is k-connected (in fact, this is the only case where edges that are in $\mathscr{G}(\boldsymbol{X}; r_k)$ but not in $\mathscr{G}(\boldsymbol{X}; r_l)$ might be needed).

Consider first an arbitrary complete graph K_{k+j}, of which Lemma 8 can be used to obtain a full hamiltonian decomposition, together with a transversal containing one edge from each of the Hamilton cycles. Now choose $k/2$ of the Hamilton cycles in the decomposition, and let T be the matching consisting of the edges of the transversal that lie in the chosen cycles.

Next, we can identify the set of vertices of K_{k+j} with the set $J \cup J''$, so that the vertices incident with edges in T are identified with the vertices of J''. Each of the original Hamilton cycles in the decomposition turns into a linear forest when restricted to the edges in G, since at least the edge in the matching T is missing in G ($J'' \subseteq J'$ has no internal edges in G). By construction, these $k/2$ forests have all the required properties.

Alternative proof of Claim 2. The proof is analogous to the one of Claim 1, so we just sketch the main differences. As before, we take a set J'' of k points in J' which are common neighbours of all points in J with respect to $\mathscr{G}(\boldsymbol{X}; r_k)$, and consider an arbitrary complete graph K_{k+j}.

If $k+j$ is even, we use the hamiltonian decomposition given by Lemma 8 to pick $(k-1)/2$ Hamilton cycles plus a perfect matching, all of them parwise edge-disjoint. Moreover, let T be a matching containing one edge from each of the Hamilton cycles. Then, we identify the set of vertices of K_{k+j} with the set $J \cup J''$, so that all vertices incident with edges in T are identified with vertices in J'', and everything follows as for Claim 1.

If $k+j$ is odd (and in particular $j > 1$), we pick $(k+1)/2$ edge-disjoint Hamilton cycles and a matching T containing one edge from each of the Hamilton cycles. Let e be the edge in T with colour $(k+1)/2$, and let u be one of its endpoints. We remove colours from $(k+j+1)/2$ of the edges in the Hamilton cycle of colour $(k+1)/2$ in a way that only a matching of all points but u remains in that colour. Finally, we identify the set of vertices of K_{k+j} with the set $J \cup J''$, so that u and the vertices incident with edges in $T \setminus \{e\}$ are identified with the vertices in J''.

6 General dimension

We can be extend our main result to general dimension, in the following sense. Modify the definition of $\boldsymbol{X} = (X_1, \ldots, X_n)$ by assuming that the X_i are chosen independently and u.a.r. from $[0,1]^d$, for some fixed integer $d \geq 2$. Redefine $\mathscr{G}(\boldsymbol{X}; r)$ analogously, using some fixed ℓ_p norm of $[0,1]^d$. Then Theorem 1 still holds for the random graph process $\big(\mathscr{G}(\boldsymbol{X}; r)\big)_{0 \leq r < \infty}$.

In fact, most of the argument in the paper is independent of d, so we shall simply sketch the main differences of those parts that need to be changed.

First of all, the definition of constant m in Section 2 is extended to

$$m = \begin{cases} 2^{d-2}(k+2-d-2/d) & \text{if } 1 \leq k < d, \\ 2^{d-1}(k+1-d-1/d) & \text{if } k \geq d \end{cases}$$

and let θ denote the volume of the unit d-dimensional ball with respect to the ℓ_p norm. Then, Proposition 3 remains valid if we change (1) to

$$\theta n {r_k}^d - \frac{2^{d-1}}{d} \log n - m \log\log n,$$

since it still follows from Theorem 8.4 in [7]. In view of that, we replace the condition on r by $\theta n r^d = \frac{2^{d-1}}{d}\log n + m \log\log n + O(1)$ in the statements of Lemmas 4 and 9 and by $\theta n r^d \sim \frac{2^{d-1}}{d}\log n$ in Lemmas 5, 6 and 7.

In order to extend the proof of Lemma 4 to general dimension d, we classify bad configurations into types according to their positition with respect to the boundary of $[0,1]^d$. More precisely, for each $i \in \{0,\dots,d\}$, let T_i denote the number of bad configurations such that X_{v_1} is at distance at most r from exactly $d-i$ *facets* (i.e. $(d-1)$-dimensional faces) of $[0,1]^d$. Setting τ_i to be the distance between X_{v_1} and the corresponding facet, we obtain by an analogous argument that, for some constant $c>0$,

$$\begin{aligned}
\mathbf{E}T_i &= O(n^{k+s+1}) \int_0^r \cdots \int_0^r \int_0^{\sqrt{d}\eta r} \rho^{d(j-2)+d-1} r^{d(k+1-j)} (\rho r^{d-1})^s \frac{e^{-cr^{d-1}n(\rho+\tau_1+\cdots+\tau_{d-i})}}{(n^{2^{d-1}/d}\log^m n)^{1/2^{d-i}}} d\rho d\tau_1 \cdots d\tau_{d-i} \\
&= O\left(\frac{(nr^d)^{k-dj+i} nr^{d-i}}{n^{2^{i-1}/d}\log^{m/2^{d-i}} n}\right) \int_0^{nr^d} \cdots \int_0^{nr^d} \int_0^{\sqrt{d}\eta n r^d} x^{d(j-1)-1+s} e^{-c(x+y_1+\cdots+y_{d-i})} dx dy_1 \cdots dy_{d-i} \\
&= O\left(\frac{(\log n)^{k-dj+i-m/2^{d-i}+(d-i)/d}}{n^{(2^{i-1}-i)/d}}\right).
\end{aligned}$$

This last expression is trivially $o(1)$ for all $i \notin \{1,2\}$, since then we have $(2^{i-1}-i)/d > 0$. To check the cases $i=1$ and $i=2$, observe that an equivalent definition of m is

$$m = \max\{2^{d-2}(k+2-d-2/d), 2^{d-1}(k+1-d-1/d)\}.$$

Hence,

$$\begin{aligned}
\mathbf{E}T_1 &= O((\log n)^{k+1-d-1/d-m/2^{d-1}+1-d(j-1)}) = O(\log^{1-d(j-1)} n) = o(1) \qquad \text{and} \\
\mathbf{E}T_2 &= O((\log n)^{k+2-d-2/d-m/2^{d-2}+1-d(j-1)}) = O(\log^{1-d(j-1)} n) = o(1).
\end{aligned}$$

The cells defined in Section 2 become d-dimensional hypercubes of side $\delta' r$, and all the remaining definitions in the that section are extended analogously. The a.a.s. events in Lemma 5 simply turn into: for each $i \in \{0,\dots,d\}$, all connected sets of cells of area at least $(1+\alpha)i\theta r^d/2^{d-1}$ touching exactly $d-i$ facets of $[0,1]^d$ contain some dense cell. The proof is completely analogous, setting $s = \lceil (1+\alpha)i\theta/(2^{d-1}\delta'^d)\rceil = \Theta(1)$, changing (2) by

$$\begin{aligned}
\sum_{t=0}^{(M-1)s} \binom{n}{t} (s\delta'^d r^d)^t (1-s\delta'^d r^d)^{n-t} &= O\left(e^{-(1+\alpha)i\theta r^d n/2^{d-1}}\right) \sum_{t=0}^{(M-1)s} (r^d n)^t \\
&= O\left(n^{-(1+\alpha)i/d+o(1)} \log^{(M-1)s} n\right),
\end{aligned}$$

and observing that there are $\Theta(1/r^i) = \Theta((n/\log n)^{i/d})$ connected sets of s cells touching $d-i$ facets of $[0,1]^d$.

Redefine *small* sets of cells to be those ones that can be embedded in a $(4d^2) \times \overset{d}{\overbrace{\cdots}} \times (4d^2)$ grid of cells (i.e. a set of cells of ℓ_∞-diametre at most $4d^2\delta' r$). Call those sets of cells that are not small *large*. In these terms, we proceed to extend Lemma 6 to general dimension $d \geq 2$. From [1] and [4], we deduce that a.a.s. $\mathscr{G}_{\mathcal{C}}[\mathcal{D}]$ has one very large component $\mathcal{D}_0$ and all (bad) components of $\mathscr{G}_{\mathcal{C}}[\mathcal{B}]$ have geometric diametre at most $r/10$ (see more specifically Lemma 5, Corollary 10 and Section 3 in [1], and also Proposition 5 and Section 3 in [4]). So in particular a.a.s. there is no bad component b such that $N(b)$ touches any pair of opposite facets of $[0,1]^d$. In order to show that a.a.s. all bad components are small it is sufficient to prove the following claim:

- A.a.s. for any large connected set of cells $\mathcal{S}$ such that $N(\mathcal{S})$ does not touch any two opposite facets of $[0,1]^d$, $N(\mathcal{S}) \setminus \mathcal{S}$ must contain some dense cell.

The proof of this claim is very similar to the one of the claim in the beginning of the proof of Lemma 6, and simply consists in bounding from below the area of some connected component of $N(\mathcal{S})\setminus\mathcal{S}$ by describing some disjoint suitable subsets contained in one topological component of $\cup(N(\mathcal{S}) \setminus \mathcal{S})$. We thus sketch only the main ideas, and describe the subsets of $\cup(N(\mathcal{S}) \setminus \mathcal{S})$ with the required properties.

For any $i \in \{0, \dots, d\}$, assume that $N(\mathcal{S})$ touches exactly $d - i$ facets of $[0,1]^d$, namely $F_1, \dots, F_{d-i}$, where w.l.o.g. $F_t = [0,1]^{t-1} \times \{0\} \times [0,1]^{d-t}$. For $i = 0$, the claim above is immediate from the extended version of Lemma 5, so we focus on the case $i > 0$.

First we need some geometric definitions that extend some of the objects we already used for $d = 2$. We call *d-ball sector* of centre $O \in \mathbb{R}^d$ to the intersection of the d-ball of centre O and radius $(1 - 2d\delta')r$ with one of the 2^d orthants that arise after translating the origin of coordinates to O. A d-ball sector has volume $\theta(1 - 2d\delta')^d r^d/2^d$. Given a cell in $[0,1]^d$, we associate to each of its 2^d corners the d-ball sector centred on the corner and in the orthant opposite to the cell. To simplify notation, let us denote the $(d-1)$th and dth coordinates in $[0,1]^d$ (or $\mathbb{R}^d$) as 'horizontal' and 'vertical' respectively, so that the usual two dimensional language can be applied when referring to these coordinates. A *vertical cylinder sector* in $\mathbb{R}^d$ is the cartesian product of a $(d-1)$-ball sector times $[a, b]$, and we say that this cylinder sector has height $b - a \geq 0$. Similarly, we can obtain a *horizontal cylinder sector* of length $b - a$ by permuting the last two coordinates. A vertical (horizontal) cylinder sector of height (length) $b - a$ has volume at least $(b - a)\theta(1 - 2d\delta')^{d-1}r^{d-1}/2^d$.

Now let c_1 and c_2 be respectively a topmost and a bottommost cell in $\mathcal{S}$ (possibly equal or not unique). We can find 2^{i-1} disjoint d-ball sectors above c_1 (associated to the top corners of c_1 which point towards the interior of $[0,1]^d$) which are contained in the same topological component of $\cup(N(\mathcal{S}) \setminus \mathcal{S})$. Similarly, we can find 2^{i-1} disjoint d-ball sectors below c_2 in the same component. Hence, some topological component of $\cup(N(\mathcal{S}) \setminus \mathcal{S})$ has area at least

$$2^{i-d}\theta(1 - 2d\delta')^d r^d.$$

Therefore, if $3 \leq i \leq d$ then $2^{i-d} > i/2^{d-1}$ and the claim follows by Lemma 5. Notice that so far we did not use the fact that $\mathcal{S}$ is large.

For the cases $i = 1, 2$, we need to achieve a better bound by finding some additional and disjoint subsets of $\cup(N(\mathcal{S}) \setminus \mathcal{S})$ with the required properties. Since $\mathcal{S}$ is large, we can assume w.l.o.g. that the vertical length of $\mathcal{S}$ is at least $4d^2\delta' r$. Consider first the case $i = 2$. Let c_3 and c_4 be respectively a leftmost and a rightmost cell in $\mathcal{S}$ (possibly equal or not unique). In addition to the previously described d-ball sectors, we consider a vertical cylinder sector of height $4d^2\delta' r$ to the right of c_4 and between the top and the bottom d-ball sectors. Similarly, consider a vertical cylinder sector of height $4d^2\delta' r$ to the left of c_3, a horizontal cylinder sector of length $\delta' r$ above c_1 and a horizontal cylinder sector of length $\delta' r$ below c_2. The total area is in this case at least

$$4\theta(1 - 2d\delta')^d r^d/2^d + 2(4d^2 + 1)\delta' r\theta(1 - 2d\delta')^{d-1}r^{d-1}/2^d =$$
$$2\theta r^d/2^{d-1}\left[(1 - 2d\delta')^d + (4d^2 + 1)\delta'(1 - 2d\delta')^{d-1}/2\right],$$

and the claim follows by Lemma 5 since $(1 - 2d\delta')^d + (4d^2 + 1)\delta'(1 - 2d\delta')^{d-1}/2 > 1$ for δ small enough.

Finally, suppose $i = 1$. This case is similar to the previous one. We consider the initial 2^i d-ball sectors as before plus an additional vertical cylinder sector of height $4d^2\delta' r$ to the right of c_4, a horizontal cylinder sector of length $\delta' r$ above c_1 and a horizontal cylinder sector of length $\delta' r$ below c_2. The total area is at least

$$2\theta(1 - 2d\delta')^d r^d/2^d + (4d^2 + 2)\delta' r\theta(1 - 2d\delta')^{d-1} r^{d-1}/2^d =$$
$$2\theta r^d/2^{d-1}\big[(1 - 2d\delta')^d + (4d^2 + 2)\delta'(1 - 2d\delta')^{d-1}/2\big],$$

and the claim follows by Lemma 5.

Lemma 7 can be extended effortlessly changing 16 to $4d^2$ and other constants appropriately. The remaining of the argument is independent of d or the extension requires just trivial changes.

References

[1] J. Balogh, B. Bollobás, and M. Walters. Hamilton cycles in random geometric graphs. Preprint, http://arxiv.org/abs/0905.4650

[2] B. Bollobás and A. Frieze, On matchings and Hamiltonian cyclesin random graphs. In *Random Graphs '83,* M. Karoński and A.Ruciński (Eds.), North-Holland, Amsterdam, pp. 23–46 (1985).

[3] J. Díaz, D. Mitsche, and X. Pérez. Sharp threshold for hamiltonicity of random geometric graphs. *SIAM Journal on Discrete Mathematics*, 21(1):57–65, 2007.

[4] M. Krivelevich, and T. Müller. Hamiltonicity of the random geometric graph. Preprint, http://arxiv.org/abs/0906.0071v2

[5] E. Lucas. Récreations Mathématiques. Vol II, Paris, 1892

[6] M.D. Penrose. On k-connectivity for a geometric random graph. *Random Structures & Algorithms*, 15(2):145–164, 1999.

[7] M.D. Penrose. *Random Geometric Graphs*, volume 5 of *Oxford Studies in Probability.* Oxford University Press, 2003.

[8] X. Pérez-Giménez, and N.C. Wormald. Disjoint Hamilton cycles in the random geometric graph. Preprint, http://arxiv.org/abs/0907.4459