

Espuny Díaz, Alberto

Hamiltonicity of graphs perturbed by a random geometric graph

Original published in: Journal of graph theory. - New York, NY [u.a.] : Wiley. - 103 (2023), 1, p. 12-22.

Original published: 2022-10-07

ISSN: 1097-0118

DOI: [10.1002/jgt.22901](https://doi.org/10.1002/jgt.22901)

[Visited: 2023-05-23]



This work is licensed under a [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Hamiltonicity of graphs perturbed by a random geometric graph

Alberto Espuny Díaz 

Institut für Mathematik, Technische Universität Ilmenau, Ilmenau, Germany

Correspondence

Alberto Espuny Díaz, Institut für Mathematik, Technische Universität Ilmenau, Weimarer Straße 25, 98684 Ilmenau, Germany.
Email: alberto.espuny-diaz@tu-ilmenau.de

Funding information

Carl-Zeiss-Stiftung

Abstract

We study Hamiltonicity in graphs obtained as the union of a deterministic n -vertex graph H with linear degrees and a d -dimensional random geometric graph $G^d(n, r)$, for any $d \geq 1$. We obtain an asymptotically optimal bound on the minimum r for which a.s. $H \cup G^d(n, r)$ is Hamiltonian. Our proof provides a linear time algorithm to find a Hamilton cycle in such graphs.

KEYWORDS

absorption, Hamiltonicity, random geometric graphs, randomly perturbed graphs

1 | INTRODUCTION

Randomly perturbed graphs are one of the thriving areas in the study of random combinatorial structures, and many interesting results in this field have been proved in recent years. The main general goal in this area is to study properties of graphs which are obtained as the union of a deterministic graph H (usually with a minimum degree condition) and a random graph G , particularly when H does not satisfy the property of interest and G is unlikely to.

Research in this direction sparked off with the work of Bohman, Frieze, and Martin [4], who studied Hamiltonicity (i.e., the property of containing a cycle which covers every vertex of a graph) in $H \cup G_{n,p}$, where H is an n -vertex graph with minimum degree at least αn , for some $\alpha \in (0, 1/2)$ (if $\alpha \geq 1/2$, then Dirac's theorem guarantees that H contains a Hamilton cycle), and $G_{n,p}$ is the binomial random graph where each pair of vertices forms an edge independently with probability p . They showed that, whenever $p \geq C(\alpha)/n$, asymptotically almost surely (a.a.s.) $H \cup G_{n,p}$ is Hamiltonian, improving on the threshold for Hamiltonicity in $G_{n,p}$ by a

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Journal of Graph Theory* published by Wiley Periodicals LLC.

logarithmic factor. Since then, Hamiltonicity has also been considered in randomly perturbed directed graphs [4, 19], hypergraphs [16, 19, 21] and subgraphs of the hypercube [9]. Many other properties have been considered as well (e.g., powers of Hamilton cycles [1, 6, 11, 23], F -factors [3, 7, 8, 15], spanning trees [5, 18, 20] or general bounded degree spanning graphs [6]), and in most cases significant improvements on the probability threshold have been achieved. To the best of our knowledge, all of these results consider (hyper/di)graphs perturbed by a binomial random structure, such as $G_{n,p}$, or its $G_{n,m}$ counterpart. Only very recently, Espuny Díaz and Girão [12] considered Hamiltonicity in graphs perturbed by a random regular graph.

In this paper we consider graphs perturbed by a random geometric graph. We consider (labelled) random geometric graphs in d dimensions, which are defined as follows. Let $V := \{1, \dots, n\}$, and let X_1, \dots, X_n be n independent uniform random variables on $[0, 1]^d$. Then, consider a positive real number r and let $E := \{\{i, j\} : \|X_i - X_j\| \leq r\}$, where $\|\cdot\|$ denotes the Euclidean norm. The resulting graph (V, E) is denoted by $G^d(n, r)$.

Hamiltonicity of random geometric graphs is fairly well understood. For dimension 2, Díaz, Mitsche and Pérez-Giménez [10] determined that the sharp threshold for Hamiltonicity is $r^* := (\log n / (\pi n))^{1/2}$ (which means that, for all $\epsilon > 0$, if $r \geq (1 + \epsilon)r^*$, then a.a.s. $G^2(n, r)$ is Hamiltonian, and if $r \leq (1 - \epsilon)r^*$, then a.a.s. $G^2(n, r)$ is not Hamiltonian), where \log stands for the natural logarithm. This result was later strengthened by Balogh, Bollobás, Krivelevich, Müller and Walters [2] and, independently, Müller, Pérez-Giménez and Wormald [22], who also extended their results to higher dimensions, proving that, for each positive integer $d \geq 2$, there exists a constant c_d such that the sharp threshold for Hamiltonicity in $G^d(n, r)$ is $(c_d \log n / n)^{1/d}$. (These two papers actually prove a stronger “hitting time” result.) The same statement is true for $d = 1$, which is a special case (see [24, Section 13.1] and the references therein, together with a remark in the introduction of [2]). All of these results also extend to other ℓ_p norms, $1 \leq p \leq \infty$, the only difference being the value of the constant $c_d = c(d, p)$. For more information about random geometric graphs, we recommend the works of Penrose [24, 25].

Our goal here is to prove the following result.

Theorem 1. *For every integer $d \geq 1$ and $\alpha \in (0, 1/2)$, there exists a constant C such that the following holds for any function $r = r(n)$ with $r \geq (C/n)^{1/d}$. Let H be an n -vertex graph with minimum degree at least αn . Then, a.a.s. $H \cup G^d(n, r)$ is Hamiltonian.*

This result is asymptotically best possible, up to the constant factor (which we will make no effort to optimise). The rest of the paper is organised as follows. In Section 2, we set our notation and also introduce a probabilistic tool which will be important for our proof. We then prove Theorem 1 in Section 3. Finally, in Section 4 we make several observations about our result, as well as some extensions that follow from its proof.

2 | PRELIMINARIES

2.1 | Notation

For any integer n , we write $[n] := \{i \in \mathbb{Z} : 1 \leq i \leq n\}$. If we write parameters in a hierarchy, we assume they are chosen from right to left. To be more precise, whenever we write

$0 < a \ll b \leq 1$, we mean that there exists an unspecified, nondecreasing function $f: \mathbb{R} \rightarrow \mathbb{R}$ such that the ensuing claim holds for all $0 < b \leq 1$ and for all $0 < a \leq f(b)$. This can be immediately generalised to longer hierarchies, and also to hierarchies where one parameter may depend on two or more other parameters. We say that a sequence of events $\{\mathcal{E}_i\}_{i \in \mathbb{N}}$ holds *asymptotically almost surely* (a.a.s.) if $\mathbb{P}[\mathcal{E}_i] \rightarrow 1$ as $i \rightarrow \infty$. In all asymptotic statements, we will ignore rounding issues whenever these do not affect the arguments.

Most of our graph-theoretic notation is standard. Given a graph G , we use $V(G)$ and $E(G)$ to denote its vertex set and edge set, respectively. We always consider labelled graphs, meaning that whenever we say that G is an n -vertex graph we may implicitly assume that $V(G) = [n]$. If G is a geometric graph (meaning here that each of its vertices is assigned to a position in \mathbb{R}^d for some integer d), then $V(G)$ may interchangeably be used to refer to the set of positions to which the vertices of G are assigned, and similarly the notation v may refer to a vertex or to its position. If $e = \{u, v\}$ is an edge, we usually abbreviate it as $e = uv$. Given any vertex $v \in V(G)$, we define $N_G(v) := \{u \in V(G) : uv \in E(G)\}$, and $d_G(v) := |N_G(v)|$ is its *degree*. We denote the minimum and maximum vertex degrees of G by $\delta(G)$ and $\Delta(G)$, respectively. Given a graph G and a set of vertices $A \subseteq V(G)$, we denote by $G[A]$ the graph on vertex set A whose edges are all edges of G which have both endpoints in A . A *path* P is a graph whose vertices can be labelled in such a way that $E(P) = \{v_i v_{i+1} : i \in [V(P)] - 1\}$. If the endpoints of a path (the first and last vertices in the labelling described above) are u and v , we sometimes refer to it as a (u, v) -path. Given a (u, v) -path P and a (v, w) -path P' such that $V(P) \cap V(P') = \{v\}$, we write PP' to denote the path obtained by concatenating P and P' (formally, this is the union graph of P and P'). If P' is a single edge vw , we will write this as Pw . Multiple concatenations will be written in the same way.

2.2 | Azuma's inequality

Let Ω be an arbitrary set (in our case, we will take $\Omega = [0, 1]^d$), and let $f: \Omega^n \rightarrow \mathbb{R}$ be some function. We say that f is L -Lipschitz, for some positive $L \in \mathbb{R}$, if, for all $x, y \in \Omega^n$ such that x and y are identical in all but one coordinate, we have that $|f(x) - f(y)| \leq L$. The following consequence of Azuma's inequality (see, e.g., the book of Janson, Łuczak, and Ruciński [17, Corollary 2.27]) will be useful for bounding the deviations of certain random variables.

Lemma 2. *Let X_1, \dots, X_n be independent random variables taking values in a set Ω , and let $f: \Omega^n \rightarrow \mathbb{R}$ be an L -Lipschitz function. Then, for any $t \geq 0$, the random variable $X := f(X_1, \dots, X_n)$ satisfies that*

$$\mathbb{P}[X \geq \mathbb{E}[X] + t] \leq e^{-\frac{t^2}{2L^2n}} \quad \text{and} \quad \mathbb{P}[X \leq \mathbb{E}[X] - t] \leq e^{-\frac{t^2}{2L^2n}}.$$

3 | PROOF OF THEOREM 1

Let $0 < 1/n_0 \ll 1/C \ll 1/d, \alpha \leq 1$, where $d \in \mathbb{N}$. Throughout, we assume that $n \geq n_0$ is a natural number. Let $r := (C/n)^{1/d}, s := 1/\lceil 2\sqrt{d}/r \rceil$ and $K := s^d n$. Our choice of C ensures that K is sufficiently large for all the ensuing claims to hold. (On an intuitive level, we may replace the hierarchy above by $0 < 1/n \ll 1/K \ll 1/d, \alpha \leq 1$, and may think of K as essentially

interchangeable with C .) Tessellate $[0, 1]^d$ using d -dimensional hypercubes of side s (intuitively, s is close to $r/(2\sqrt{d})$, but possibly slightly smaller to guarantee that the tessellation above exists). We refer to each of the smaller d -dimensional hypercubes as a *cell*, and denote the set of all cells as \mathcal{C} . Observe that $|\mathcal{C}| = s^{-d} = n/K$. We say that two cells $c_1, c_2 \in \mathcal{C}$ are *friends* if their boundaries intersect (in particular, if the boundaries share a single point, they count as intersecting). It follows that each cell is friends with at most $3^d - 1$ other cells. Given any set of points $S \subseteq [0, 1]^d$, we say that a cell is *sparse* with respect to S if it contains at most $R := 2 \cdot 3^d$ points, and we call it *dense* with respect to S otherwise.

Consider a labelling of the vertices of $V(H)$ as v_1, \dots, v_n , and let X_1, \dots, X_n be independent uniform random variables on $[0, 1]^d$. We may assume that none of the variables take their values in the boundary of any of the cells, as this is an event that occurs with probability 1. Consider the random geometric graph $G = G^d(n, r)$ on the vertex set of H obtained by assigning position X_i to v_i . Note that, by the definition of s , if some $v \in V(H)$ lies in a cell $c \in \mathcal{C}$, then it is joined by an edge of G to all other vertices in c , as well as to all vertices in cells which are friends of c .

Let us provide here a brief sketch of the proof. When considering the random geometric graph $G^d(n, r)$, one can show that a.s. “most” cells will be dense with respect to the vertex set, and only a small proportion of them will be sparse. Now remove all sparse cells from consideration and define an auxiliary graph Γ , whose vertex set is the set of remaining cells, and where two cells are joined by an edge if they are friends. A simple algorithm allows us to construct a cycle containing all the vertices which lie in cells that form a connected component in Γ . Indeed, fix one such component and consider a spanning subtree T . By definition, T has bounded degree. One can then “walk” along the edges of T to visit all cells, using each edge of T exactly twice, and end in the starting cell. While doing this, and since the edges of T correspond to cells which are friends, one can incorporate the vertices in the cells into a path (where a suitable number of vertices must be chosen each time a cell is visited). Eventually, this path gets closed into a cycle when the walk returns to the starting cell. The bound on the maximum degree of T and the definition of sparse cells (with respect to $V(H)$) are crucial for this process to work. This simple idea has been at the heart of many results about Hamiltonicity in random geometric graphs [2, 10, 22]. Since “most” cells are dense and we can apply the previous algorithm to each component of Γ , we can obtain a set of cycles covering almost all vertices. To obtain a Hamilton cycle, all that remains is to join the cycles into a single, longer cycle and to incorporate the missing vertices. It is for this purpose that we will need to use the edges of H (and also be somewhat careful when choosing the cycles above). The following definition will be crucial to achieve our goal.

Given any pair of (not necessarily distinct) vertices $u, v \in V(H)$, we say that a cell is $\{u, v\}$ -*linked* (with respect to H) if it contains two distinct vertices $x, y \in V(H) \setminus \{u, v\}$ such that $ux, vy \in E(H)$; if it does not contain such a pair, we say that it is $\{u, v\}$ -*unlinked*. In the particular case when $u = v$, we will say that the cell is v -linked/unlinked to mean that it is $\{v, v\}$ -linked/unlinked. One can easily see an intuitive reason why this notion is useful. Say that a cell c is sparse with respect to $V(H)$ and contains two vertices u and v , and say we find a $\{u, v\}$ -linked cell c' which is dense with respect to $V(H)$. Then, when “walking” through the component of Γ containing c' , we may actually leave this component to reach u , incorporate all vertices in c into the path we are constructing, and then go back from v to c' . This will allow us to incorporate all vertices in sparse cells into the cycles we constructed above. A similar approach will allow us to join the different cycles.

We now proceed to formally prove Theorem 1. We begin with the following claim.

Claim 1. The following properties hold a.s.:

- (i) The number of cells which are sparse with respect to $V(H)$ is at most $e^{-K/2n}$.
- (ii) For each pair of (not necessarily distinct) vertices $u, v \in V(H)$, the number of $\{u, v\}$ -unlinked cells is at most $4e^{-\alpha K/4n}$.

Proof.

- (i) Let $f : ([0, 1]^d)^n \rightarrow \mathbb{Z}_{\geq 0}$ be a function which, given a set S of points $x_1, \dots, x_n \in [0, 1]^d$, returns the number of cells which are sparse with respect to S . Note that f is a 1-Lipschitz function.

For each $v \in V(H)$ and each cell c , we have that $\mathbb{P}[v \in c] = s^d$. Thus, since the variables X_i are independent, for a fixed cell c we have that

$$\begin{aligned} \mathbb{P}[c \text{ is sparse with respect to } V(H)] &= \sum_{i=0}^R \binom{n}{i} s^{di} (1 - s^d)^{n-i} \\ &\leq (R + 1) \binom{n}{R} s^{dR} (1 - s^d)^{n-R} \leq 2R \frac{K^R}{R!} \left(1 - \frac{K}{n}\right)^n \leq K^R e^{-K}. \end{aligned}$$

(Here the first and third inequalities rely on the hierarchy we established at the beginning of the proof; in particular, we may think of K as being much larger than R . We also use the fact that $s^d = K/n$.) Let $Y := f(X_1, \dots, X_n)$ be the number of cells which are sparse with respect to $V(H)$, so $\mathbb{E}[Y] \leq K^{R-1} e^{-K} n$. Since f is 1-Lipschitz, it follows by Lemma 2 that

$$\mathbb{P}[Y \geq e^{-K/2} n] \leq \mathbb{P}[Y \geq 2K^{R-1} e^{-K} n] \leq e^{-\Theta(n)}.$$

- (ii) We proceed in a similar way. For each positive integer m , let $g_m : ([0, 1]^d)^m \rightarrow \mathbb{Z}_{\geq 0}$ be a function which, given a set S of m points $x_1, \dots, x_m \in [0, 1]^d$, returns the number of cells which contain at most one point. Clearly, g_m is 1-Lipschitz for every $m \in \mathbb{N}$. Given any set $V \subseteq V(H)$, we say that a cell is V -vacant if it contains at most one of the vertices of V .

Fix two (not necessarily distinct) vertices $u, v \in V(H)$, and let Z be the number of $\{u, v\}$ -unlinked cells. Let $N := N_H(u) \cap N_H(v)$. We split the analysis into two cases. Assume first that $\ell := |N| \geq \alpha n/2$. Let i_1, \dots, i_ℓ be the indices of the vertices of H which lie in N . It follows that, for a fixed cell c ,

$$\mathbb{P}[c \text{ is } N\text{-vacant}] = (1 - s^d)^\ell + \ell s^d (1 - s^d)^{\ell-1} \leq K e^{-\alpha K/4}.$$

Let $Z' := g_\ell(X_{i_1}, \dots, X_{i_\ell})$, so $\mathbb{E}[Z'] \leq e^{-\alpha K/4} n$. Note that every $\{u, v\}$ -unlinked cell is N -vacant, so $Z \leq Z'$. Since g_ℓ is 1-Lipschitz, by Lemma 2 we conclude that

$$\mathbb{P}[Z \geq 2e^{-\alpha K/4} n] \leq e^{-\Theta(n)}.$$

Assume now that $\ell < \alpha n/2$ and let $S := N_H(u) \setminus (N_H(v) \cup \{v\})$ and $T := N_H(v) \setminus (N_H(u) \cup \{u\})$. Note that $|S|, |T| \geq \alpha n/2$. By following the same argument as above, we have that, with probability at least $1 - e^{-\Theta(n)}$, the number of S -vacant cells is at most $2e^{-\alpha K/4} n$. The same holds for the number of T -vacant cells. Note that every $\{u, v\}$ -unlinked cell must be S -vacant or T -vacant. We thus conclude that

$$\mathbb{P}[Z \geq 4e^{-\alpha K/4}n] \leq e^{-\Theta(n)}.$$

Finally, the statement holds by a union bound over all pairs of vertices $\{u, v\} \subseteq V(H)$. □

Condition on the event that G satisfies the properties of the statement of Claim 1, which holds a.a.s. Let \mathcal{C}_s be the set of cells which are sparse with respect to $V(H)$ (to which we simply refer as “sparse cells” from now on), and let $\mathcal{C}_d := \mathcal{C} \setminus \mathcal{C}_s$. We define an auxiliary graph Γ with vertex set \mathcal{C}_d where two cells are joined by an edge whenever they are friends. In particular, $\Delta(\Gamma) \leq 3^d - 1$.

Claim 2. The number of connected components of Γ is at most $e^{-K/3}n$.

Proof. By Claim 1(i), there are at most $e^{-K/2}n$ sparse cells. Recall that the total number of cells is n/K . Thus, the number of components of Γ of size at least $n^{1/d}/(4K^{1/d})$ is at most $4(n/K)^{1-1/d} \leq e^{-K/3}n/2$.

Now consider any component γ of Γ with fewer than $n^{1/d}/(4K^{1/d})$ cells. We claim that the number of sparse cells which are friends with some cell of γ is at least $2^d - 1$. Indeed, choose an arbitrary cell $c_0 \in V(\gamma)$, and choose a corner of the hypercube $[0, 1]^d$ which is at distance at least $1/3$ from c_0 in each coordinate direction. Let us assume, without loss of generality, that said corner is $\{1\}^d$. Now, for $i \geq 0$, we follow an iterative process. Consider the cube C_i of side length $2s$ having c_i in a corner and growing in all directions towards $\{1\}^d$; this is a union of 2^d cells, all of which are friends with each other. If C_i only contains sparse cells (other than c_i), we are done (and stop the iterative process), so assume otherwise. Then, there must be some dense cell $c_{i+1} \subseteq C_i \setminus c_i$. This new cell c_{i+1} , being a friend of c_i , also lies in γ , and is a translation of c_i by a vector all whose coordinates are nonnegative and bounded from above by s . This process keeps going until either we find an index i such that C_i contains $2^d - 1$ sparse cells or we reach one of the sides of the hypercube $[0, 1]^d$. However, reaching the side requires at least $1/(3s) = n^{1/d}/(3K^{1/d})$ iterations, which contradicts the assumption on the order of γ , so our claim must indeed hold.

On the other hand, trivially no sparse cell can be friends with more than $3^d - 1$ cells which lie in distinct components of Γ . Now, a simple double counting argument guarantees that the number of components of Γ with fewer than $n^{1/d}/(4K^{1/d})$ cells is at most $2^d e^{-K/2}n \leq e^{-K/3}n/2$. Together with the bound in the first paragraph, this completes the proof. □

With the statements of Claims 1 and 2, we can already begin the main part of the proof. We begin by setting up some notation. Let $\mathcal{C}^0 \subseteq \mathcal{C}_s$ be the set of all cells which contain no vertices of $V(H)$, and let $\mathcal{C}^1 \subseteq \mathcal{C}_s$ be the set of all cells which contain exactly one vertex of $V(H)$. Let $\mathcal{C}^2 := \mathcal{C}_s \setminus (\mathcal{C}^0 \cup \mathcal{C}^1)$, and $\mathcal{C}^* := \mathcal{C}^1 \cup \mathcal{C}^2$. Recall that, for any $v \in V(H)$, we use v to refer both to the vertex and to its position in $[0, 1]^d$. Given any cell $c \in \mathcal{C}_d$, let $\Gamma(c)$ be the connected component of Γ which contains c . Let $\mathcal{F} := \mathcal{C}_s$ and $\mathcal{F}^* := \{c \in \mathcal{C}_d : |V(\Gamma(c))| = 1\}$. By Claim 1(i) and Claim 2 we have that $|\mathcal{F} \cup \mathcal{F}^*| \leq 2e^{-K/3}n$. Both \mathcal{F} and \mathcal{F}^* constitute sets of “forbidden” cells which we will avoid when connecting vertices from different cells via edges of H . These sets will be updated as we choose edges of H to construct the Hamilton cycle. Indeed, as we choose edges of H , each cell which contains a vertex incident to any of these edges will be

added to \mathcal{F} , so that the different choices of edges do not interact with each other. These edges of H are chosen “from a cell”, and \mathcal{F}^* will contain all dense cells from which we must choose edges in a “correct” way (namely, ensuring that our process works), and thus will be avoided when choosing edges of H from other cells. We claim that, for the rest of the proof, we will always have $\mathcal{C}_s \subseteq \mathcal{F}$ and

$$|\mathcal{F} \cup \mathcal{F}^*| \leq e^{-K/6}n, \quad (1)$$

and assume so throughout. This bound will follow from the fact that we update \mathcal{F} and \mathcal{F}^* at most $2e^{-K/3}n$ times, and each time the size of their union will increase by at most 3.

Consider an auxiliary graph $\Gamma' := \Gamma$. We are first going to modify this graph Γ' into a connected graph. We will update Γ' in $t - 1$ steps, where $t \leq e^{-K/3}n$ is the number of components of Γ (see Claim 2). In each of these steps, we will add exactly one edge to Γ' , connecting two of its components. This auxiliary edge will correspond to a way in which we will later connect the cycles which we will construct in each component; we build the structure necessary for this at the same time as we update Γ' . Our definitions of \mathcal{F} and \mathcal{F}^* are crucial in guaranteeing that the upcoming process can be carried out. In particular, \mathcal{F} and \mathcal{F}^* will always be disjoint, none of the components of (the current form of) Γ' will be contained in \mathcal{F} , and \mathcal{F}^* will always contain at most one cell of each component of Γ' . Given any cell $c \in \mathcal{C}_d$, let $\Gamma'(c)$ denote the connected component of Γ' which contains c . Initialise a set of vertices F_d and two sets of edges E_d and E_d^* as empty sets. We proceed as follows.

1. For each $i \in [t - 1]$, choose a smallest component γ of Γ' , and choose an arbitrary cell $c \in V(\gamma) \setminus \mathcal{F}$ (which exists since $V(\gamma) \not\subseteq \mathcal{F}$). Let u_c and v_c be two arbitrary distinct vertices in c . Choose an arbitrary $\{u_c, v_c\}$ -linked cell $c'(c) \in \mathcal{C} \setminus (\mathcal{F} \cup \mathcal{F}^* \cup V(\gamma))$; its existence follows from Claim 1(ii), (1) and the fact that γ is a smallest component of Γ' . Add the edge $\{c, c'(c)\}$ to Γ' . Let $x_c \in N_H(u_c) \cap c'(c)$ and $y_c \in N_H(v_c) \cap c'(c)$ be two distinct vertices. Add u_c, v_c, x_c and y_c to F_d , add $u_c x_c$ and $v_c y_c$ to E_d^* , and add $u_c v_c$ and $x_c y_c$ to E_d . Then, add c and $c'(c)$ to \mathcal{F} (if $c \in \mathcal{F}^*$, remove it from this set, so that \mathcal{F} and \mathcal{F}^* remain disjoint). Finally, if $|V(\Gamma'(c)) \setminus \mathcal{F}| = 1$, add this remaining cell to \mathcal{F}^* .

Each iteration of step 1 reduces the number of components by one, so it follows that, after we perform all iterations, Γ' is connected. Moreover, a moment of thought reveals that \mathcal{F}^* must now be empty.

We next define some *absorbing paths* which will be used to incorporate all vertices in sparse cells into a Hamilton cycle. We define these iteratively in $|\mathcal{C}^*|$ steps. We proceed as follows.

2. For each $c \in \mathcal{C}^1$, let v_c be the vertex contained in c , and choose an arbitrary v_c -linked cell $c'(c) \in \mathcal{C} \setminus \mathcal{F}$; note that such a cell exists by Claim 1(ii) and (1). Choose two distinct vertices $x_c, y_c \in N_H(v_c) \cap c'(c)$. Note that $e_c := x_c y_c \in E(G)$, and define $P_c := x_c v_c y_c$. Then, add $c'(c)$ to \mathcal{F} .
3. For each $c \in \mathcal{C}^2$, let u_c and v_c be two arbitrary distinct vertices in c and choose an arbitrary $\{u_c, v_c\}$ -linked cell $c'(c) \in \mathcal{C} \setminus \mathcal{F}$, which again exists by Claim 1(ii) and (1). Let $x_c \in N_H(u_c) \cap c'(c)$ and $y_c \in N_H(v_c) \cap c'(c)$ be two distinct vertices, and let P'_c be any (u_c, v_c) -path which contains all vertices in c (recall that such a path exists because $G[V(H) \cap c]$ is a complete graph). Let $e_c := x_c y_c \in E(G)$ and $P_c := x_c P'_c y_c$. Then, add $c'(c)$ to \mathcal{F} .

Once this process is finished, let $F_s := \{x_c, y_c : c \in \mathcal{C}^*\}$ and $E_s := \{e_c : c \in \mathcal{C}^*\}$.

Let $F := F_s \cup F_d$ and $E := E_s \cup E_d$. Note that, by construction, for each $c \in \mathcal{C}$ we have that $|F \cap c| \leq 2$ and $|E(G[V(H) \cap c]) \cap E| \leq 1$. We are now ready to construct the Hamilton cycle. The main step for this is to construct a cycle in each component γ of Γ . We make sure that these cycles contain all edges of E spanned by the vertices in the cells of the corresponding component. For each component γ of Γ , we proceed as follows.

- Let T be a spanning tree of γ . In particular, $\Delta(T) < 3^d$. Consider an arbitrary traversal of T which, starting at a given cell, goes through every edge of T twice and ends in the starting cell (this can be given, for example, by a DFS on T taking any cell c_0 as a root). This traversal takes $m := 2(|V(\gamma)| - 1)$ steps, each step corresponding to an edge of T . We use this traversal to construct a cycle $\mathcal{C}(\gamma)$ as follows.

Assume the traversal starts in a given cell c_0 . Choose a vertex $x_0 \in (V(H) \cap c_0) \setminus F$ and let $P_0 := x_0$; this will be the beginning of a path which we will grow into $\mathcal{C}(\gamma)$. For notational purposes, set $V(P_{-1}) := \emptyset$. For each $i \in [m]$ we define a path P_i as follows. Let c be the current cell in our traversal, and let $x_{i-1} \in (V(H) \cap c) \setminus F$ be the last vertex of P_{i-1} . Let c' be the next cell of the traversal. Because c and c' are friends, every vertex in c' is joined to every vertex in c by an edge of G . Choose an arbitrary vertex $x_i \in (V(H) \cap c') \setminus (F \cup V(P_{i-1}))$. If this is the last time that c is visited in the traversal of T , let P'_i be any path with vertex set $(V(H) \cap c) \setminus V(P_{i-2})$ having x_{i-1} as an endpoint and such that, if the vertices in c span some edge $e \in E$, then $e \in E(P'_i)$, and let $P_i := P_{i-1}P'_ix_i$; otherwise, simply let $P_i := P_{i-1}x_i$. To complete the cycle, let P' be an (x_0, x_m) -path whose internal vertices are all vertices of $(V(H) \cap c_0) \setminus V(P_m)$ and such that, if the vertices in c_0 span some edge $e \in E$, then $e \in E(P')$. We then set $\mathcal{C}(\gamma) := P_m \cup P'$. Observe that every cell contains at least $R = 2 \cdot 3^d$ vertices and is visited at most 3^d times throughout the traversal; this, together with the fact that no cell spans more than one of the edges of E , guarantees that the choices of vertices described throughout the process can always be carried out.

Let \mathcal{C} denote the graph which is the union of all the cycles constructed in step 4. In particular, $E \subseteq E(\mathcal{C})$. We can now combine the cycles into a single cycle spanning all vertices in cells of \mathcal{C}_d by letting $\mathcal{C}' := (\mathcal{C} \setminus E_d) \cup E_d^*$. To complete the proof, for each $c \in \mathcal{C}^*$, replace the edge $e_c \in E(\mathcal{C}')$ by P_c . □

Remark. By using results from percolation theory, for $d \geq 2$ one can show that a.s. the graph Γ actually contains a “giant” component which contains, say, more than 9/10 of the cells. This fact can be used instead of Claim 2 to streamline our proof, simplifying some of its technical details, such as the need for the set \mathcal{F}^* . Indeed, step 1 can be avoided entirely: for every cell c which does not lie in this giant component we may find a suitable $\{u, v\}$ -linked cell which does, for some $u, v \in V(H) \cap c$, as in steps 2 and 3. Step 4 can then be applied only to this giant component. This approach, however, does not work when $d = 1$, as Γ will a.s. not contain any component of linear size.

4 | FINAL REMARKS

Theorem 1 is asymptotically best possible in the sense that, for each $\alpha \in (0, 1/2)$, if we let $r = (c/n)^{1/d}$ for a sufficiently small constant c , then there exist graphs H with minimum degree αn such that $H \cup G^d(n, r)$ is a.s. not Hamiltonian. Indeed, let H be a complete unbalanced

bipartite graph with parts A and B of sizes αn and $(1 - \alpha)n$, respectively. Clearly, if $G^d(n, r)[B]$ contains more than αn isolated vertices, then $H \cup G^d(n, r)$ cannot contain a Hamilton cycle. Our claim thus follows immediately from the following lemma (where we also allow α to depend on n).

Lemma 3. *Let $d \geq 1$ be an integer, and let V_d be the volume of the ball of radius 1 in d dimensions. If $1/2 > \alpha = \alpha(n) = \omega(n^{-1/2})$, then a.a.s. $G^d(n, (-\log(3\alpha/2)/(2V_d n))^{1/d})$ contains at least $4\alpha n/3$ isolated vertices.*

Proof. Let $f: ([0, 1]^d)^n \rightarrow \mathbb{Z}$ be a function that, given a set of n points $x_1, \dots, x_n \in [0, 1]^d$, returns the number of isolated vertices in the geometric graph constructed on these points with radius $(-\log(3\alpha/2)/(2V_d n))^{1/d}$. One can readily check that f is an L -Lipschitz function, for some $L = L(d)$.

Fix an arbitrary $i \in [n]$. Let X_1, \dots, X_n be independent uniform random variables on $[0, 1]^d$. We are first going to compute the probability that i is isolated in $G = G^d(n, (-\log(3\alpha/2)/(2V_d n))^{1/d})$. This is the probability that no X_j with $j \neq i$ is assigned to a position within distance r of X_i , that is,

$$\mathbb{P}[i \text{ is isolated}] \geq (1 + \log(3\alpha/2)/(2n))^{n-1} \geq e^{\log(3\alpha/2)} = 3\alpha/2.$$

Let $X := f(X_1, \dots, X_n)$ denote the number of isolated vertices in G , so $\mathbb{E}[X] \geq 3\alpha n/2$. Since f is an L -Lipschitz function, by Lemma 2 we conclude that

$$\mathbb{P}[X \leq 4\alpha n/3] \leq \mathbb{P}[X \leq \mathbb{E}[X] - \alpha n/6] \leq e^{-\Theta(\alpha^2 n)} = o(1),$$

where the last equality relies on the lower bound on α . □

A natural question arises from Theorem 1: What can we say about Hamiltonicity of randomly perturbed graphs when α is allowed to be a function of n (which tends to 0)? This same question was recently considered by Hahn-Klimroth, Maesaka, Mogge, Mohr and Parczyk [14] when the random graph G is binomial, and they extended the result of Bohman, Frieze and Martin [4] by proving that $G = G(n, \Theta(-\log \alpha/n))$ is both sufficient and necessary. In view of Lemma 3, the following seems a natural conjecture (and, if true, it would be best possible up to the value of the constant C).

Conjecture 4. *For every integer $d \geq 1$, there exists a constant C such that the following holds. Let $\alpha = \alpha(n) \in (0, 1/2)$, and let H be an n -vertex graph with minimum degree at least αn . Then, a.a.s. $H \cup G^d(n, (-C \log \alpha/n)^{1/d})$ is Hamiltonian.*

We also wish to remark upon two features of our proof of Theorem 1. First, the proof is constructive, meaning that it provides an algorithm to find Hamilton cycles in $H \cup G^d(n, r)$. In particular, if the properties of Claim 1 hold (which occurs a.a.s.), it provides a deterministic algorithm that outputs a Hamilton cycle in $H \cup G^d(n, r)$. Furthermore, observe that, throughout the proof, we actually do not need Claim 1(ii) to hold for every pair of vertices, but only for those that we pick throughout the process, which are only linearly many. This means that the properties of Claim 1 can be checked in $\mathcal{O}(n^2)$ time, which is linear in the size of

$H \cup G^d(n, r)$. Then, the construction of the Hamilton cycle also takes $\mathcal{O}(n^2)$ time. This follows directly from the proof, and can be checked by retracing the steps.

Second, all throughout the paper we have considered the ℓ_2 Euclidean norm for simplicity. Our proof generalises directly to all ℓ_p norms, $1 \leq p \leq \infty$, by adjusting some of the constants.

Furthermore, we note that, under the same conditions as in the statement of Theorem 1, we can actually show that $H \cup G^d(n, r)$ is pancyclic. Indeed, our proof can easily be modified for this. From the Hamilton cycle that we construct, given that it contains many subpaths whose vertices actually form cliques in $G^d(n, r)$, one can iteratively reduce the number of vertices in the cycle. This can be balanced with also removing some of the paths which correspond to sparse cells, as well as leaves of the auxiliary graph Γ' , to prove that cycles of all lengths can be constructed.

Finally, we remark that our result opens the door to questions about other spanning structures in dense graphs perturbed by random geometric graphs. In this direction, Espuny Díaz and Hyde [13] very recently generalised Theorem 1 to deal with powers of Hamilton cycles, showing that, for k independent of n , the required radius for the a.a.s. containment of the k th power of a Hamilton cycle in this setting is the same (up to constant factors depending on k) as that required for Hamiltonicity. This has a number of direct consequences, such as for questions related to F -factors or 2-universality (which directly implies pancyclicity).

ACKNOWLEDGEMENTS

I would like to thank Xavier Pérez-Giménez for some very helpful discussions about the topic of this paper. This study has been partially supported by the Carl Zeiss Foundation. Open Access funding enabled and organized by Projekt DEAL.

ORCID

Alberto Espuny Díaz  <http://orcid.org/0000-0001-5066-3820>

REFERENCES

1. S. Antoniuk, A. Dudek, C. Reiher, A. Ruciński, and M. Schacht, *High powers of Hamiltonian cycles in randomly augmented graphs*, *J. Graph Theory*. **98** (2021), no. 2, 255–284. <https://doi.org/10.1002/jgt.22691>
2. J. Balogh, B. Bollobás, M. Krivelevich, T. Müller, and M. Walters, *Hamilton cycles in random geometric graphs*, *Ann. Appl. Probab.* **21** (2011), 1053–1072. <https://doi.org/10.1214/10-AAP718>
3. J. Balogh, A. Treglown, and A. Z. Wagner, *Tilings in randomly perturbed dense graphs*, *Combin. Probab. Comput.* **28** (2019), 159–176. <https://doi.org/10.1017/S0963548318000366>
4. T. Bohman, A. Frieze, and R. Martin, *How many random edges make a dense graph Hamiltonian?* *Random Structures Algorithms*. **22** (2003), 33–42. <https://doi.org/10.1002/rsa.10070>
5. J. Böttcher, J. Han, Y. Kohayakawa, R. Montgomery, O. Parczyk, and Y. Person, *Universality for bounded degree spanning trees in randomly perturbed graphs*, *Random Structures Algorithms*. **55** (2019), 854–864. <https://doi.org/10.1002/rsa.20850>
6. J. Böttcher, R. Montgomery, O. Parczyk, and Y. Person, *Embedding spanning bounded degree graphs in randomly perturbed graphs*, *Mathematika*. **66** (2020), 422–447. <https://doi.org/10.1112/mtk.12005>
7. J. Böttcher, O. Parczyk, A. Sgueglia, and J. Skokan, *Triangles in randomly perturbed graphs*, *Combin. Probab. Comput.* (2022), 1–31. <https://doi.org/10.1017/S0963548322000153>
8. J. Böttcher, O. Parczyk, A. Sgueglia, and J. Skokan, *Cycle factors in randomly perturbed graphs*, *Procedia Comput. Sci.* **195** (2021), 404–411. *Proceedings of the XI Latin and American Algorithms, Graphs and Optimization Symposium*. <https://doi.org/10.1016/j.procs.2021.11.049>
9. P. Condon, A. Espuny Díaz, A. Girão, D. Kühn, and D. Osthus, *Hamiltonicity of random subgraphs of the hypercube*, *Mem. Am. Math. Soc.* https://www.ams.org/cgi-bin/mstrack/accepted_papers/memo

10. J. Díaz, D. Mitsche, and X. Pérez, *Sharp threshold for Hamiltonicity of random geometric graphs*, SIAM J. Discrete Math. **21** (2007), 57–65. <https://doi.org/10.1137/060665300>
11. A. Dudek, C. Reiher, A. Ruciński, and M. Schacht, *Powers of Hamiltonian cycles in randomly augmented graphs*, Random Structures Algorithms. **56** (2020), 122–141. <https://doi.org/10.1002/rsa.20870>
12. A. Espuny Díaz and A. Girão, *Hamiltonicity of graphs perturbed by a random regular graph*, Random Struct. Algor.
13. A. Espuny Díaz and J. Hyde, *Powers of Hamilton cycles in dense graphs perturbed by a random geometric graph*, arXiv e-prints (2022). arXiv: 2205.08971.
14. M. Hahn-Klimroth, G. S. Maesaka, Y. Mogge, S. Mohr, and O. Parczyk, *Random perturbation of sparse graphs*, Electron. J. Combin. **28** (2021), no. 2, 12. Research paper p. 2.26 <https://doi.org/10.37236/9510>
15. J. Han, P. Morris, and A. Treglown, *Tilings in randomly perturbed graphs: Bridging the gap between Hajnal–Szemerédi and Johansson–Kahn–Vu*, Random Structures Algorithms. **58** (2021), 480–516. <https://doi.org/10.1002/rsa.20981>
16. J. Han and Y. Zhao, *Hamiltonicity in randomly perturbed hypergraphs*, J. Combin. Theory Ser. B. **144** (2020), 14–31. <https://doi.org/10.1016/j.jctb.2019.12.005>
17. S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000. <https://doi.org/10.1002/9781118032718>
18. F. Joos and J. Kim, *Spanning trees in randomly perturbed graphs*, Random Structures Algorithms. **56** (2020), 169–219. <https://doi.org/10.1002/rsa.20886>
19. M. Krivelevich, M. Kwan, and B. Sudakov, *Cycles and matchings in randomly perturbed digraphs and hypergraphs*, Combin. Probab. Comput. **25** (2016), 909–927. <https://doi.org/10.1017/S0963548316000079>
20. M. Krivelevich, M. Kwan, and B. Sudakov, *Bounded-degree spanning trees in randomly perturbed graphs*, SIAM J. Discrete Math. **31** (2017), 155–171. <https://doi.org/10.1137/15M1032910>
21. A. McDowell and R. Mycroft, *Hamilton ℓ -cycles in randomly perturbed hypergraphs*, Electron. J. Combin. **25** (2018), 30. Paper No. 4.36. <https://doi.org/10.37236/7671>
22. T. Müller, X. Pérez-Giménez, and N. Wormald, *Disjoint Hamilton cycles in the random geometric graph*, J. Graph Theory. **68** (2011), 299–322. <https://doi.org/10.1002/jgt.20560>
23. R. Nenadov and M. Trujić, *Sprinkling a few random edges doubles the power*, SIAM J. Discrete Math. **35** (2021), no. 2, 988–1004. <https://doi.org/10.1137/19M125412X>
24. M. Penrose, *Random geometric graphs*, Oxf. Stud. Probab., vol. 5, Oxford University Press, Oxford, 2003.
25. M. Penrose, *Lectures on random geometric graphs*, Random Graphs, Geometry and Asymptotic Structure (N. Fountoulakis and D. Hefetz, eds.), Cambridge University Press, Cambridge, 2016, pp. 67–101. <https://doi.org/10.1017/CBO9781316479988.004>

How to cite this article: A. Espuny Díaz, *Hamiltonicity of graphs perturbed by a random geometric graph*, J. Graph Theory. 2023;103:12–22. <https://doi.org/10.1002/jgt.22901>