On the treewidth of random geometric graphs and percolated grids

Anshui Li* Tobias Müller

August 18, 2015

Abstract

In this paper, we study the treewidth of the random geometric graph, obtained by dropping n points onto the square $[0, \sqrt{n}]^2$ and connect pairs of points by an edge if their distance is at most $r = r(n)$. We prove a conjecture of Mitsche and Perarnau [19] stating that, with probability going to one as $n \to \infty$, the treewidth the random geometric graph is $\Theta(r\sqrt{n})$ when $\liminf r > r_c$, where r_c is the threshold radius for the appearance of the giant component. The proof makes use of a comparison to standard bond percolation and with a little bit of extra work we are also able to show that, with probability tending to one as $k \to \infty$, the treewidth of graph we get by retaining each edge of the $k \times k$-grid with probability p is $\Theta(k)$ if $p > 1/2$ and $\Theta(\sqrt{\log k})$ if $p < 1/2$.

Keywords: Random geometric graphs, treewidth.
MSC classification (2010): 05C80, 60D05, 82B43.

1 Introduction and main results

The *random geometric graph* $\mathcal{G}(n, r)$ is the random graph obtained by taking n points $X_1, \ldots, X_n$ i.i.d. uniformly at random from the square $[0, \sqrt{n}]^2$, and joining X_i and X_j by an edge if their euclidean distance is at most r. Here $r = r(n)$ may and often does depend on n. To avoid having to deal with annoying trivial special cases we assume $r \leq \sqrt{2n}$ throughout the paper. The study of random geometric graphs essentially goes back to Gilbert [7] who defined a very similar model in 1961. For this reason random geometric graphs are often also called the *Gilbert model* of random graphs. Random geometric graphs have been the subject of a considerable research effort in the last two decades. As a result, detailed information is now known on various aspects such as (k-)connectivity [22, 23], the largest component [24], the chromatic number and clique number [20, 17], the (non-)existence of Hamilton cycles [2, 21], monotone graph properties in general [8] and the simple random walk on the graph [4]. One of the most well-known phenomena in random geometric graphs is the "sudden emergence of a giant component". By this we mean that there exists critical value r_c such that if $\limsup r < r_c$ then, a.a.s., every component of $G(n, r)$ has $O(\log n)$ vertices, whereas if $\liminf r > r_c$ then, a.a.s, there exists a "giant" component of with $\Omega(n)$ vertices. Here and in the rest of the paper, we say that a sequence of events $(E_n)_n$ holds *asymptotically almost surely* (abbreviation: a.a.s.)

*The first author acknowledges support from CSC-UU scholarship grant.

if $\lim_{n\to\infty} \mathbb{P}(E_n) = 1$. The exact value of r_c is not known at this time, but simulations suggest that the exact value is approximately 1.2 (see [24]). For more details, and proofs, on the giant component phenomenon and background on random geometric graphs in general we refer the reader to the monograph [24].

In the present paper, we consider the *treewidth* of random geometric graphs. The treewidth of a graph was introduced by Halin in [10] and independently but later by Robertson and Seymour in [25]. It is a graph parameter that in a sense measures how similar a given graph is to a tree. (We postpone the precise – and technical – definition of treewidth until the next section of the paper, in order to streamline the introduction.) Treewidth plays an important role in modern algorithmic graph theory. Many NP-hard algorithmic decision problems have for instance been shown to be polynomially solvable when restricted to the class of instances with a bounded tree-width. In fact, a striking result of Courcelle [5] states that any algorithmic decision problem that can be expressed in monadic second order logic, can be solved in linear time for the class of graphs with bounded treewidth. An example of a decision problem that is NP-hard in general and can be expressed in monadic second order is k-colourability (for any fixed k). As random geometric graphs have been used extensively as models for modeling communication networks, this motivated Mitsche and Perarnau [19] to consider the treewidth of random geometric graphs. They proved that if $r \in (0, r_c)$ is fixed then, a.a.s., $\mathrm{tw}(G(n,r)) = \Theta(\log n/\log\log n)$, while if $r > C$ where C is a large constant then, a.a.s., $\mathrm{tw}(G(n,r)) = \Theta(r\sqrt{n})$. Mitsche and Perarnau [19] also conjectured that the second result should extend all the way to the critical value. Here we will prove their conjecture:

Theorem 1.1 *If $r = r(n)$ is such that $\liminf r > r_c$, where r_c is the critical value for the emergence of the giant component, then, a.a.s. as $n \to \infty$, $\mathrm{tw}(G(n,r)) = \Theta(r\sqrt{n})$.*

Our proof of Theorem 1.1 makes use of a comparison to bond percolation on $\mathbb{Z}^2$. Recall that this refers to the infinite random graph obtained by retaining each edge of the familiar integer lattice with probability p and discarding it with probability $1-p$, independently of the choices for all other edges. We will denote by $\Gamma(k,p)$ the restriction of this process to the $k \times k$ integer grid. I.e., $\Gamma(k,p)$ has vertex set $[k]^2$ and for every pair of points $u, v \in [k]^2$ with euclidean distance equal to one, we add an edge with probability p, independently of the choices for all other pairs. (Here and in the rest of the paper we use the notation $[k] := \{1, \dots, k\}$.) For the proof of Theorem 1.1 we only need to consider the treewidth of $\Gamma(k,p)$ when p is very close to one, but with a little bit of extra work we are able to obtain the following result in addition to Theorem 1.1.

Theorem 1.2 *If $p \in (0,1)$ is fixed then, a.a.s. as $k \to \infty$:*

$$\mathrm{tw}\left(\Gamma(k,p)\right) = \begin{cases} \Theta(k) & \text{if } p > 1/2\ ; \\ \Theta(\sqrt{\log k}) & \text{if } p < 1/2\ . \end{cases}$$

Note that k is the *square root* of the number of vertices of $\Gamma(k,p)$.

2 Notation and preliminaries

In this section, we give some definitions and results which we will need in the sequel. We start with the precise definition of treewidth. For a graph $G = (V,E)$ on n vertices, we call $(T, \mathcal{W})$ a *tree decomposition* of G, where $\mathcal{W}$ is a set of vertex subsets $W_1, ..., W_s \subset V$, called *bags*, and T is a forest with vertices in $\mathcal{W}$, such that:

(i) $\cup_{i=1}^{s} W_i = V$;

(ii) For any $e = uv \in E$ there exists a set $W_i \in \mathcal{W}$ such that $u, v \in W_i$;

(iii) For any $v \in V$, the subgraph induced by the $W_i \ni v$ is connected as a subgraph of T.

The *width* of a tree-decomposition is $w(T, \mathcal{W}) = \max_{1 \leq i \leq s} \mid W_i \mid -1$, and the *treewidth* of a graph G can be defined as

$$\mathrm{tw}(G) := \min_{(T,\mathcal{W})} w(T, \mathcal{W}),$$

where the minimum is of course taken over all tree decompositions $(T, \mathcal{W})$ of G. From the definition of treewidth, one can see that the treewidth of a tree is one, while the treewidth of a k-clique is $k-1$. Let us also observe that if H is a subgraph of G, then $\mathrm{tw}(H) \leq \mathrm{tw}(G)$, and if G is a graph with connected components $G_1, ..., G_m$, then $\mathrm{tw}(G) = \max_{1 \leq i \leq m} \mathrm{tw}(G_i)$.

Given an edge xy of graph G, the graph G/xy is obtained from G by *contracting* the edge xy. That is, to obtain G/xy, we identify the vertices x and y and remove all resulting loops and duplicate edges. A graph H is a *minor* of G if it is a subgraph of graph obtained from G by a sequence of edge-contractions. Again, one can see from the definitions that if H is a *minor* of G, $\mathrm{tw}(H) \leq \mathrm{tw}(G)$.

Alon, Seymour and Thomas [1] proved the following powerful result, bounding the treewidth of graphs without a given minor.

Theorem 2.1 ([1]) *If G does not have H as a minor, then* $\mathrm{tw}(G) \leq |V(H)|^{\frac{3}{2}} \cdot \sqrt{|V(G)|}$.

In this paper we will make use of the following immediate corollary:

Corollary 2.2 *There exists a constant $C > 0$ such that every planar graph G satisfies* $\mathrm{tw}(G) \leq C\sqrt{|V(G)|}$.

Throughout the paper we will denote by $\Gamma(k)$ $(:= \Gamma(k, 1))$ the $k \times k$ grid. The next observation appears as Exercise 16 in Chapter 12 of [6].

Lemma 2.3 $\mathrm{tw}(\Gamma(k)) = k$.

For one of our lower bounds on the treewidth, we will need the following lemma which links the treewidth of a graph and the existence of a partition of its vertex set with special properties. A vertex partition $V = \{A, S, B\}$ is a *balanced k-partition* if $|S| = k + 1$, there are no edges in G between a vertex in A and a vertex in B, and $\frac{1}{3}(n-k-1) \leq |A|, |B| \leq \frac{2}{3}(n-k-1)$. In this case, S is called a *balanced separator*. The following result connecting balanced partitions and treewidth is due to Kloks [14], which provides a necessary condition for a graph to have a treewidth of certain size.

Lemma 2.4 ([14]) *Let G be a graph and suppose that $tw(G) \leq k \leq |V(G)| - 1$. Then G has a balanced k-partition.*

We say that $A \subseteq \{0, 1\}^n$ is an up-set if whenever we take a point of A and we change one of its coordinates into a one, then the resulting point is still in A. We will use the following lemma later on.

Lemma 2.5 (Harris' lemma, [11]) *Let $A, B \subseteq \{0,1\}^n$ be up-sets and let $X = (X_1, \dots, X_n)$ be a vector of independent Bernoulli random variables. Then $\mathbb{P}(X \in A \cap B) \geq \mathbb{P}(X \in A)\mathbb{P}(X \in B)$.*

By a slight abuse of notation, throughout this paper we will denote the graph with vertex set $\mathbb{Z}^2$ and an edge $vw \in E(\mathbb{Z}^2)$ if and only if $\|v - w\| = 1$ also by $\mathbb{Z}^2$. Recall that bond percolation on $\mathbb{Z}^2$ refers to the random process where we keep each edge of $\mathbb{Z}^2$ with probability p and discard it with probability $1-p$, independently of all other edges. The edges that are kept are referred to as *open* and the discarded edges as *closed*. If $R := \{a, \dots, b\} \times \{c, \dots, d\}$ is an axis-parallel rectangle, then we say that R has a *horizontal crossing* if there is an open path that stays inside R and connects the left side $\{a\} \times \{c, \dots, d\}$ to the right side $\{b\} \times \{c, \dots, d\}$. A vertical crossing is defined similarly. We denote by $H(R)$ the event that there is a horizontal crossing of R, and $V(R)$ the event that there is a vertical crossing of R. In the sequel, we will use the following well-known result on bond percolation on $\mathbb{Z}^2$ with $p > \frac{1}{2}$. A proof can for instance be found in [3] (Lemma 8 on page 64).

Lemma 2.6 *If $p > 1/2$ then $\lim_{k\to\infty} \mathbb{P}(H([3k] \times [k])) = 1$.*

In words, when $p > 1/2$ then the probability of crossing a $3k \times k$ rectangle in the long direction can be made arbitrarily close to one by making k large.

Formally speaking, we can describe bond percolation on $\mathbb{Z}^2$ as a random vector X taking values in $\{0,1\}^{E(\mathbb{Z}^2)}$. Here $X_e = 1$ if e is open, and $X_e = 0$ otherwise. In the standard setup, the coordinates X_e are i.i.d. Bernoulli random variables. Of course one can also consider more general bond percolation models in which the coordinates are not independent. We say that such a bond percolation model Y is 1*-independent* if, for every pair of sets $S, T \subseteq E(\mathbb{Z}^2)$ with the property that no edge in S shares an endpoint with any edge in T, the random vectors $(Y_e)_{e\in S}$ and $(Y_e)_{e\in T}$ are independent. Recall that a coupling of two random objects X, Y is a joint probability space on which both are defined (and have the correct marginal distributions). The following result is a reformulation of a special case of a result by Liggett, Schonmann and Stacey [15]:

Theorem 2.7 ([15]) *There exists a function $\pi : [0,1] \to [0,1]$ such that, $\lim_{p\uparrow 1} \pi(p) = 1$, and the following holds. Suppose that Y follows a 1-independent bond percolation model on $\mathbb{Z}^2$ and set $p := \inf_{e\in E(\mathbb{Z}^2)} \mathbb{P}(Y_e = 1)$. Then there exists a coupling of Y with a standard (i.e. independent) bond percolation X with $\mathbb{P}(X_e = 1) = \pi(p)$, such that, almost surely, $X_e \leq Y_e$ for all $e \in E(\mathbb{Z}^2)$.*

In words, this last theorem says that every 1-independent bond percolation model contains the edges of an independent bond percolation model, and the edge probability $\pi(p)$ of this independent bond percolation approaches one as $p := \inf_{e\in E(\mathbb{Z}^2)} \mathbb{P}(Y_e = 1)$ approaches one.

When working with random geometric graphs, it is often useful to consider a *Poissonized* version of the random geometric graph. We define $G_{\text{Po}}(n, r)$ analogously to $G(n, r)$ except that now we we take the points of a Poisson process of intensity one on $[0, \sqrt{n}]^2$ and then build our graph on that as before. Equivalently, we can say that we throw $N_n \stackrel{d}{=} \text{Po}(n)$ i.i.d. uniform points onto $[0, \sqrt{n}]$ and then build the graph on those as before. Working with the Poissonized version is often useful in proofs because of the convenient independence properties of the Poisson process. Recall that if $N_n \stackrel{d}{=} \text{Po}(n)$ then $\mathbb{P}(N_n > (1+\varepsilon)n) = o(1)$, as can for instance be seen by Chebyschev's inequality. Using a straightforward coupling and rescaling, this gives the following observation:

Corollary 2.8 *There is a coupling such that for every $r = r(n)$, a.a.s.,$G_{Po}\left((1-\varepsilon)n, r\sqrt{1-\varepsilon}\right)$ is a subgraph of $G(n,r)$.*

It of course also makes sense to simply consider the random geometric graph built on a Poisson process $\mathcal{P}$ of intensity one on all of the plane $\mathbb{R}^2$. This is the well-known *continuum percolation* model defined originally by Gilbert [7]. We remark that Gilbert and several other sources in the literature fix $r = 1$ and allow the intensity of the Poisson process to vary, but it is easily seen that this is equivalent to the setting where we vary r and the intensity of the Poisson process is fixed to be one. Note that $G_{\mathrm{Po}}(n,r)$ is just the restriction of continuum percolation to the square $[0,\sqrt{n}]^2$. Let us also remark that the critical r_c for the "emergence of a giant component" in $G(n,r)$ is the same as the critical value for the existence of an infinite component in continuum percolation (see [24], Chapters 9 and 10). Similarly to the case of bond percolation on $\mathbb{Z}^2$, we can define crossing events for continuum percolation. Our definition follows Meester and Roy [18]. For $R = [a,b]\times[c,d] \subseteq \mathbb{R}^2$ an axis-parallel rectangle, we say that $H(R)$ holds (i.e. there is a horizontal crossing of R) if it is possible to draw a continuous curve between the right and left side that stays inside R and is completely covered by the balls of radius $r/2$ centered on the points of $\mathcal{P}$. Note that this in particular implies that there is a path between a vertex that is within $r/2$ of the left side of R, and a vertex within $r/2$ of the right side of R such that all other vertices of the path are either inside R or within distance $r/2$ of R. We have the following analogue of Lemma 2.6.

Lemma 2.9 ([18], Corollary 4.1) *If $r > r_c$ then $\lim_{a\to\infty}\mathbb{P}(H([0,3a]\times[0,a])) = 1$.*

We say that an event A defined with respect to the Poisson process $\mathcal{P}$ is *increasing* if whenever A is true for some set of points $X = \{x_1, x_2, \dots\} \subseteq \mathbb{R}^2$ (i.e. some realization of $\mathcal{P}$), then A also holds for any set X' that contains X. We have the following analogue of Lemma 2.5 above.

Lemma 2.10 ([18], Theorem 2.2) *If A, B are increasing events (wrt. $\mathcal{P}$) then $\mathbb{P}(A\cap B) \geq \mathbb{P}(A)\mathbb{P}(B)$.*

3 The treewidth of the percolated grid $\Gamma(k,p)$

3.1 When p is large

Instead of proving the $p > 1/2$ part of Theorem 1.2 directly, we first prove the following weaker version:

Proposition 3.1 *There exist constants $c > 0$ and $p < 1$ such that $\mathrm{tw}(\Gamma(k,p)) \geq ck$ a.a.s.*

If $A \subseteq \mathbb{Z}^2$ is finite and connected (as a subgraph of $\mathbb{Z}^2$) then there is a well defined "surrounding cycle" $\mathrm{surr}(A)$ in the dual lattice $(\mathbb{Z}^2)^* = \mathbb{Z}^2 + (\frac{1}{2},\frac{1}{2})$ (that separates A from infinity, and every other cycle in $(\mathbb{Z}^2)^*$ that separates A from infinity contains $\mathrm{surr}(A)$ in its interior). For $A \subseteq [k]^2$ connected, we define $\mathrm{outer}(A)$ to be the set of edges of $\Gamma(k)$ that cross $\mathrm{surr}(A)$. (See Figure 1 for a depiction.)

We will make use of the following straightforward observation. We include a proof for completeness.

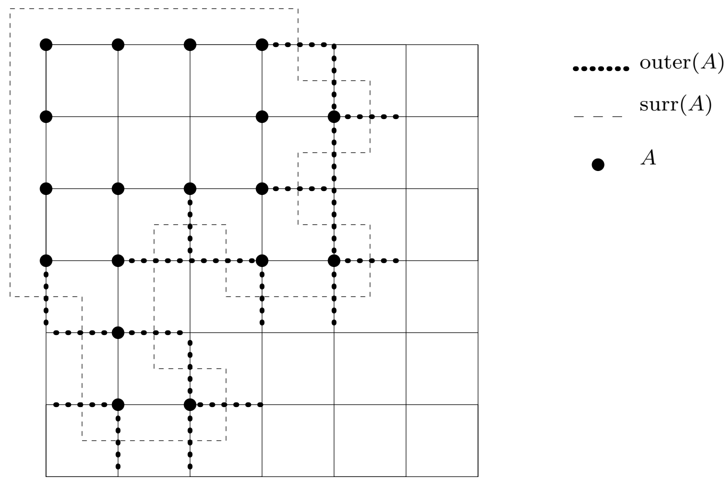

Figure 1: Depiction of surr(A) and outer(A) for a set $A \subseteq [7]^2$.

Lemma 3.2 *Suppose that $A \subseteq [k]^2$ is connected (as a subgraph of $\Gamma(k)$) and does not contain a horizontal crossing of $[k]^2$. Then $|\operatorname{outer}(A)| \geq \max(\sqrt{|A|}, \operatorname{surr}(A)/4)$.*

Proof: First suppose that A contains a vertical crossing. Since A does not contain a horizontal crossing, surr(A) must contain a (dual) path that separates the left edge of $[k]^2$ from its right edge. This implies that outer(A) contains at least k edges. Hence outer(A) $\geq \sqrt{|A|}$. Note that, since A can intersect at most one of the vertical sides of $[k]^2$, we have that the number of edges of surr(A) that do not intersect edges of $\Gamma(k)$ is at most $|\operatorname{surr}(A)| - |\operatorname{outer}(A)| \leq 3k$. This shows $|\operatorname{outer}(A)| \geq |\operatorname{surr}(A)|/4$.

Let us then assume that A contains neither a horizontal nor a vertical crossing. Let $a := |\pi_x(A)|, b := |\pi_y(A)|$, where π_x, resp. π_y, denote the projection onto the x-axis, resp. the y-axis. Clearly we have that $|A| \leq a \cdot b$ and $|\operatorname{outer}(A)| \geq a + b$. Thus,

$$\sqrt{|A|} \leq \max(a, b) \leq a + b \leq |\operatorname{outer}(A)|.$$

Also note that $|\operatorname{surr}(A)| - |\operatorname{outer}(A)| \leq a + b$. So certainly $|\operatorname{outer}(A)| \geq |\operatorname{surr}(A)|/4$. ■

We say a set $A \subseteq [k]^2$ is *dirty* (wrt. $\Gamma(k, p)$) if **1)** A is connected (as a subgraph of $\Gamma(k)$), **2)** A intersects at most three sides of $[k]^2$, and **3)** at most half of the edges of outer(A) are open in $\Gamma(k, p)$. We say that a vertex $v \in [k]^2$ is dirty if it is contained in some dirty set.

Lemma 3.3 *There exists a $p_0 < 1$ such that whenever $p \geq p_0$ then, a.a.s., there are at most $k^2/10^{10}$ dirty vertices.*

Proof: Let Y denote the number of vertices contained in some dirty set A with $|\operatorname{surr}(A)| \geq k^{.01}$ and let Z denote the number of vertices contained in some dirty set A with $|\operatorname{surr}(A)| < k^{.01}$. It is easy to see that the number of cycles in $(\mathbb{Z}^2)^*$ that have length ℓ and that surround a given vertex $v \in \mathbb{Z}^2$ is at most $\ell 4^\ell$. This allows us to bound the expectation of Y as follows:

$$
\begin{aligned}
\mathbb{E}(Y) &\leq k^2 \sum_{\ell \geq k^{0.01}} 4^\ell \binom{\ell}{\ell/8} (1-p)^{\ell/8} \\
&\leq k^2 \sum_{\ell \geq k^{0.01}} 8^\ell (1-p)^{\ell/8} \\
&= \frac{k^2 \left(8(1-p)^{1/8}\right)^{k^{0.01}}}{1-8(1-p)^{1/8}} \\
&= o(1),
\end{aligned}
$$

where we have used that $|\operatorname{outer}(A)| \geq |\operatorname{surr}(A)|/8$ in the first line, that $\binom{\ell}{\ell/8} \leq 2^\ell$ in the second line, and where the last equality holds provided p_0 was chosen sufficiently close to one (and $p \geq p_0$). In particular, for p_0 sufficiently close to one and $p > p_0$, we have that $Y = 0$ a.a.s.

Next, we consider Z. For $v \in [k]^2$, we denote by E_v the event that v is contained in a dirty A with $|\operatorname{surr}(A)| < k^{.01}$. We have

$$
\begin{aligned}
\mathbb{P}(E_v) &\leq \sum_{\ell \leq k^{0.01}} \ell 4^\ell \binom{\ell}{\ell/8} (1-p)^{\ell/8} \\
&\leq \sum_{\ell \leq k^{0.01}} \ell \left(8(1-p)^{1/8}\right)^\ell \\
&\leq \frac{8(1-p)^{1/8}}{\left(1-8(1-p)^{1/8}\right)^2}. \\
&\leq 10^{-11},
\end{aligned}
$$

where the last inequality holds provided p_0 is chosen sufficiently close to one (and $p \geq p_0$).

On the other hand, we clearly have

$$\mathbb{P}(E_v) \geq (1-p)^4.$$

Hence, we have that $k^2(1-p)^4 \leq \mathbb{E}Z = \sum_v \mathbb{P}(E_v) \leq k^2/10^{11}$. In particular, $\mathbb{E}Z = \Theta(k^2)$.

Next we consider the second moment of Z. Observe that if $|u-v|_\infty \geq 3 \cdot k^{.01}$, then E_u and E_v are independent. (Here $|(x,y)|_\infty = \max(|x|,|y|)$ denotes the familiar L_∞-norm.) This allows us to write

$$
\begin{aligned}
\mathbb{E}Z^2 &= \sum_{u,v} \mathbb{P}(E_u \cap E_v) \\
&= \sum_v \mathbb{P}(E_v) \sum_{|u-v|_\infty < 3k^{.01}} \mathbb{P}(E_u|E_v) + \sum_v \mathbb{P}(E_v) \sum_{|u-v|_\infty \geq 3k^{.01}} \mathbb{P}(E_u|E_v) \\
&\leq \sum_v \mathbb{P}(E_v) 36k^{.02} + \sum_{u,v} \mathbb{P}(E_v)\mathbb{P}(E_u) \\
&= \mathbb{E}Z \cdot o(k^2) + (\mathbb{E}Z)^2 \\
&= (1+o(1))\,(\mathbb{E}Z)^2\,.
\end{aligned}
$$

This shows that $\operatorname{Var}(Z) = o\left((\mathbb{E}Z)^2\right)$. An application of Chebyschef's inequality shows:

$$\begin{aligned}
\mathbb{P}(Z > k^2/10^{10}) &\le \mathbb{P}(|Z - \mathbb{E}Z| \ge \tfrac{9}{10}\mathbb{E}Z) \\
&\le (\tfrac{10}{9})^2 \cdot \tfrac{\mathrm{Var}(Z)}{(\mathbb{E}Z)^2} \\
&= o(1).
\end{aligned}$$

In conclusion, we have seen that, when p_0 is sufficiently close to one and $p_0 < p \le 1$, then $Y = 0$ a.a.s. and $Z \le k^2/10^{10}$ a.a.s., which obviously implies the lemma. ■

Proof of Proposition 3.1: Let us pick $1 > p > p_0$ with p_0 as provided by Lemma 3.3. Then, a.a.s., $\Gamma(k,p)$ has no more than $k^2/10^{10}$ dirty vertices. In the remainder of the proof we therefore assume we are given a subgraph $G \subseteq \Gamma(k)$ for which there are at most $k^2/10^{10}$ dirty vertices, but which is otherwise arbitrary. We will show that any such G satisfies $\mathrm{tw}(G) \ge k/1000$.

Aiming for a contradiction, we assume that there exists some balanced partition $\{A, S, B\}$ of $V(G) = [k]^2$ with $|S| < k/1000$.

We first observe that we can assume without loss of generality that A does not contain a horizontal crossing. For, if it does then B cannot contain a vertical crossing (otherwise A, B would not be disjoint). Hence, by applying symmetry (switching the roles of A, B and rotating by 90 degrees) we can indeed assume A does not contain a horizontal crossing. Observe that

$$|A| \ge k^2 - |B| - |S| \ge k^2 - 2k^2/3 - k/1000 \ge k^2/10.$$

Let $A_1, \dots, A_m$ denote the connected components of A (connected when considered as subgraphs of G). Let us set

$$\mathcal{I} := \{i : A_i \text{ is not dirty}\}, \quad A' := \bigcup_{i \in \mathcal{I}} A_i.$$

Since the total number of dirty vertices is less that $k^2/10^{10}$, we have

$$|A'| \ge |A| - k^2/10^{10} \ge k^2/100.$$

Note that every edge (in G) between a vertex of A' and a vertex of $[k^2] \setminus A'$ must in fact connect a vertex of A' to a vertex of S. Hence, it follows that

$$\begin{aligned}
4|S| &\ge \textstyle\sum_{i \in \mathcal{I}} |\,\mathrm{outer}(A_i)|/2 \\
&\ge \textstyle\sum_{i \in \mathcal{I}} \sqrt{|A_i|}/2 \\
&\ge \sqrt{|A'|}/2 \\
&\ge k/20.
\end{aligned}$$

(Here we have used Lemma 3.2 for the second line and the concavity of the square root function for the third line.) So it follows that $k/1000 \ge |S| \ge k/80$, a contradiction!

This shows that there is no balanced partition with $|S| < k/1000$, which implies that $\mathrm{tw}(G) \ge k/1000$ by Kloks' lemma (Lemma 2.4). ■

3.2 When $p > 1/2$

We are now ready to prove the first part of Theorem 1.2 with the help of Proposition 3.1.

Proof of Theorem 1.2, the case $p > 1/2$: Our proof is an application of a standard technique for comparing supercritical percolation to percolation with p close to one, by means of Lemma 2.6 and Theorem 2.7. See for instance [3], pages 74–75.

Let p_0 be as provided by Proposition 3.1, and let π be as provided by Theorem 2.7. We now pick p_1 such that $\pi(p_1) > p_0$. By Lemma 2.6, we can find an $a \in \mathbb{N}$ such that $\mathbb{P}(H([3a] \times [a])) > \sqrt[3]{p_1}$.

For R a $3a \times a$ rectangle, we define the event $E(R) := H(R) \cap V(R_L) \cap V(R_R)$ where R_L denotes the leftmost $a \times a$ subrectangle, and R_R denotes the rightmost $a \times a$ rectangle (see Figure 13 on page 74 of [3] for a depiction). If R is a $a \times 3a$ rectangle the we define $E(R) := V(R) \cap H(R_B) \cap H(R_T)$ with R_B, resp. R_T, the bottom, resp. top, $a \times a$ subrectangle of A. Note that, by choice of a and Harris' lemma, we have that $\mathbb{P}(E(R)) > p_1$ for every $3a \times a$ or $a \times 3a$ rectangle R.

We now define a (dependent) bond percolation model Y on $\mathbb{Z}^2$ as follows. We declare the horizontal edge between (i, j) and $(i + 1, j)$ open in Y if $E(\{2ai + 1, \dots, 2ai + 3a\} \times \{2aj + 1, \dots, 2aj + a\})$ holds; and similarly the edge between (i, j) and $(i, j + 1)$ is open in Y if $E(\{2ai + 1, \dots, 2ai + a\} \times \{2aj + 1, \dots, 2aj + 3a\})$ holds. It is not difficult to see that Y is in fact 1-independent. Hence, by Theorem 2.7, $Y \geq X$, where X is standard (independent) percolation on $\mathbb{Z}^2$ with edge-probability $> p_0$.

We can view $\Gamma(k, p)$ as the restriction of the (independent, edge-probability p) percolation process to the $k \times k$ grid $[k]^2$. We let Γ_X, resp. Γ_Y, denote the subgraph that X, resp. Y, defines on $[\ell]^2$, where $\ell := \lfloor k/2a \rfloor$. Note that we have chosen ℓ so that each of the rectangles corresponding to the edges of Γ_Y is contained in $[k]^2$. Observe that by construction (and Proposition 3.1) we have that, a.a.s.:

$$\mathrm{tw}(\Gamma_Y) \geq \mathrm{tw}(\Gamma_X) = \Omega(\ell) = \Omega(k).$$

Next, we remark that Γ_Y is in fact a minor of $\Gamma(k, p)$ (under the natural coupling associated with the construction of Y). To see this, we can proceed as follows: If $E(R)$ holds with R a $3a \times a$ rectangle that corresponds to some edge of Γ_Y, then we do a sequence of contractions that will identify all vertices of R_L that participate in (horizontal or vertical) crossings of R_L into a single vertex x, we produce a vertex y via contractions in R_R similarly, and then we contract the remaining edges of a long, horizontal crossing of R into a single edge that connects x and y. If we do this for each rectangle corresponding to an edge of Γ_Y, then discard any unneeded vertices (making sure to keep exactly one vertex in each $a \times a$ square that corresponds to a vertex $(i, j) \in [\ell]^2$ that was not incident to any edge of Y), then we obtain a graph isomorphic to Γ_Y.

Since Γ_Y is a minor of $\Gamma(k, p)$, we have $\mathrm{tw}(\Gamma(k, p)) \geq \mathrm{tw}(\Gamma_Y) = \Omega(k)$, a.a.s., as required. ■

3.3 When $p < 1/2$

In this section, we prove the upper and lower bound of the treewidth $\Gamma(k, p)$ for $p < 1/2$. We need the following result from percolation theory, that is originally due to Kesten [12, 13].

Theorem 3.4 ([12, 13]) *Consider bond percolation on $\mathbb{Z}^2$ and let C_0 denote the number of vertices in the cluster (component) of the origin. For each $p < 1/2$ there exists $\lambda(p) > 0$ such that*

$$\mathbb{P}(|C_0| \geq n) \leq e^{-n\lambda(p)},$$

for all $n \geq 0$.

This has the following easy consequence:

Corollary 3.5 *If* $0 < p < 1/2$ *then, a.a.s., all components of* $\Gamma(k,p)$ *have* $O(\log k)$ *vertices.*

Proof: Let us fix $0 < p < 1/2$ and let $\lambda(p)$ be as provided by Theorem 3.4. Let $K := 100/\lambda(p)$. Observe that, for every $v \in [k]^2$ and $\ell \in \mathbb{N}$, the probability that is in a component of order $\geq \ell$ in $\Gamma(k,p)$ is no more that the probability that $|C_0|$ exceeds ℓ. Thus, we can conclude:

$$\begin{aligned}
\mathbb{P}(\Gamma(k,p) \text{ has a component of size} \geq K\log k) &\leq k^2 \exp[-100\log k] \\
&= \exp[-98\log k] \\
&= o(1).
\end{aligned}$$

■

Since the treewidth of a graph equals the maximum of the treewidth of its components, and all components of $\Gamma(k,p)$ are planar, the required upper bound for $\mathrm{tw}(\Gamma(k,p))$ in the case when $p < 1/2$ follows immediately using Corollary 2.2:

Corollary 3.6 *If* $0 < p < 1/2$ *then, a.a.s.,* $\mathrm{tw}(\Gamma(k,p)) = O(\sqrt{\log k})$.

The following lemma now completes the proof of Theorem 1.2.

Lemma 3.7 *Fix* $0 < p < 1/2$ *then, a.a.s.,* $\mathrm{tw}(\Gamma(k,p)) = \Omega(\sqrt{\log k})$.

Proof: We fix a $\varepsilon = \varepsilon(p)$ (small, to be determined later), and we set $\ell := \lceil\sqrt{\varepsilon\log k}\rceil$. We now fix $N := \lfloor k/(\ell+1)\rfloor^2 = \Omega(k^2/\log k)$ (vertex-) disjoint $\ell \times \ell$-subgrids $G_1, \ldots, G_N$ in $[k]^2$. We will say that the subgrid G_i is *intact* if all of its edges are present in $\Gamma(k,p)$. By independence of the events that the G_i-s are intact, we have:

$$\begin{aligned}
\mathbb{P}(\text{at least one } G_i \text{ is intact}) &= 1-(1-p^{2\ell(\ell-1)})^N \\
&\geq 1-\exp[-Np^{2\ell(\ell-1)}] \\
&\geq 1-\exp[-Np^{\ell^2}] \\
&= 1-\exp[-Np^{\varepsilon\log k}].
\end{aligned}$$

Next, note that

$$Np^{\varepsilon\log k} = \Omega\left(\frac{k^2}{\log k}\cdot\exp[\varepsilon\log p\log k]\right) = \Omega(\exp[2\log k - \log\log k + \varepsilon\log p\log k]).$$

Hence, provided we chose $\varepsilon < -2/\log p$, we have that $Np^{\varepsilon\log k} \to \infty$ and hence also

$$\mathbb{P}(\text{at least one } G_i \text{ is intact}) = 1-o(1).$$

Hence, by Lemma 2.3, and since $\mathrm{tw}(H) \leq \mathrm{tw}(G)$ if $H \subseteq G$, we have that $\mathrm{tw}(\Gamma(k,p)) \geq \ell = \Omega(\sqrt{\log k})$ a.a.s. ■

Corollary 3.6 and Lemma 3.7 together give the $p < 1/2$ part of Theorem 1.2.

4 Proof of Theorem 1.1

Since Mitsche and Perarnau [19] have already shown the result is true when $r = r(n)$ is larger than some fixed constant C, we only need to consider the case when $r_c < \liminf r \leq \limsup r \leq C$. Note that in this case $\Theta(r\sqrt{n})$ simplifies to $\Theta(\sqrt{n})$. Moreover, by monotonicity, we see that for any such sequence r, a.a.s., $\mathrm{tw}(G(n,r)) \leq \mathrm{tw}(G(n,C)) = O(\sqrt{n})$ by Mitsche and Perarnau's result. Hence, we only need to prove an a.a.s. lower bound for the treewidth of order $\Omega(\sqrt{n})$. Using Corollary 2.8 and monotonicity, Theorem 1.1 follows if we establish:

Lemma 4.1 *For every fixed $r > r_c$, we have $\mathrm{tw}(G_{Po}(n,r)) = \Omega(\sqrt{n})$ a.a.s.*

Proof: The proof is almost exactly the same as the proof of the $p > 1/2$ case of Theorem 1.2 above. Again, we let p_0 be as provided by Proposition 3.1, we let π be as provided by Theorem 2.7, and we pick p_1 such that $\pi(p_1) > p_0$. Using Lemma 2.9, we find an a such that $\mathbb{P}(H([0,3a] \times [0,a])) > \sqrt[3]{p_1}$. For R a $3a \times a$ or $a \times 3a$ rectangle we define $E(R)$ as in the proof of the $p > 1/2$ case of Theorem 1.2 above. By choice of a and Lemma 2.10 we have that $\mathbb{P}(E(R)) > p_1$ for any such rectangle.

We again define a 1-independent bond percolation model Y on $\mathbb{Z}^2$, by declaring the horizontal edge between (i,j) and $(i+1,j)$ open in Y if $E([2ai, 2ai+3a] \times [2aj, 2aj+a])$ holds; and the edge between (i,j) and $(i,j+1)$ is open in Y if $E([2ai, 2ai+a] \times [2aj, 2aj+3a])$ holds. (Note that 1-independence holds provided we chose a sufficiently large.) Again Theorem 2.7 gives that $Y \geq X$, where X is standard (independent) percolation on $\mathbb{Z}^2$ with edge-probability $> p_0$.

We set $k := \lfloor \sqrt{n}/2a \rfloor$, and we let Γ_X, resp. Γ_Y, be the restriction of X, resp. Y, to $[k]^2$. Arguing analogously to the way we did in the proof of the $p > 1/2$ case of Theorem 1.2, we see that Γ_Y is in fact a minor of $G_{\mathrm{Po}}(n,r)$ (under the natural coupling we get from the construction of Y). Hence, using Proposition 3.1, we get that a.a.s.:

$$\mathrm{tw}(G_{\mathrm{Po}}(n,r)) \geq \mathrm{tw}(\Gamma_Y) \geq \mathrm{tw}(\Gamma_X) = \Omega(k).$$

Since $k = \Theta(\sqrt{n})$ this concludes the proof. ■

5 Discussion and further work

Together with the work of Mitsche and Perarnau [19], our Theorem 1.1 provides an almost complete picture of the behaviour of the treewidth of random geometric graphs, up to the order of the leading constants:

Corollary 5.1 *A.a.s,*

$$\mathrm{tw}(G(n,r)) = \begin{cases} \Theta\left(\frac{\log n}{\log\log n}\right) & \text{if } 0 < \liminf r \leq \limsup r < r_c, \\ \Theta\left(r\sqrt{n}\right) & \text{if } \liminf r > r_c. \end{cases}$$

Interestingly, by a result of McDiarmid [16], the clique number of the random geometric graphs is a.a.s. equal to $(1+o(1))\log n/\log\log n$ when r is constant. This gives rise to the following natural questions.

Question 5.2 *Suppose* $0 < \liminf r \leq \limsup r < r_c$.
Is $\text{tw}(G(n,r)) = (1+o(1))\log n/\log\log n$ *a.a.s.?*
Is $\text{tw}(G(n,r)) = \omega(G(n,r))$ *a.a.s.?*

Of course we would also be very interested to learn the precise leading constants for the supercritical case. With our methods and those of Mitsche and Perarnau [19] the following natural conjecture still seems out of reach:

Conjecture 5.3 *Suppose that* $r > r_c$ *is fixed. Then there exists a* $c = c(r)$ *such that* $\text{tw}(G(n,r)) = (c+o(1))\sqrt{n}$ *a.a.s.*

Another tantalizing question is what happens precisely at the critical point. Based on widely believed conjectures on the "critical exponents" for two-dimensional percolation (see [9], Chapters 9 and 10), we offer the following conjectures:

Conjecture 5.4 *A.a.s.,* $\text{tw}(G(n,r_c)) = n^{\frac{91}{192}+o(1)}$.

Conjecture 5.5 *A.a.s.,* $\text{tw}(\Gamma(k,1/2)) = k^{\frac{91}{96}+o(1)}$.

We have made two separate conjectures and added some slack in the exponent so that there is a bit more hope that the at least one of the conjectures will be solved.

References

[1] N. Alon, P. Seymour, and R. Thomas. A separator theorem for nonplanar graphs. *J. Amer. Math. Soc.*, 3(4):801–808, 1990.

[2] J. Balogh, B. Bollobás, M. Krivelevich, T. Müller, and M. Walters. Hamilton cycles in random geometric graphs. *Annals of Applied Probability*, 21(3):1053–1072, 2011.

[3] B. Bollobas and O Riordan. *Percolation*. Cambridge University Press, 2006.

[4] C. Cooper and A. Frieze. The cover time of random geometric graphs. *Random Structures Algorithms*, 38(3):324–349, 2011.

[5] B. Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inform. and Comput.*, 85(1):12–75, 1990.

[6] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, fourth edition, 2010.

[7] E. N. Gilbert. Random plane networks. *J. Soc. Indust. Appl. Math.*, 9:533–543, 1961.

[8] A. Goel, S. Rai, and B. Krishnamachari. Monotone properties of random geometric graphs have sharp thresholds. *Ann. Appl. Probab.*, 15(4):2535–2552, 2005.

[9] G. Grimmett. *Percolation*, volume 321 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, second edition, 1999.

[10] R Halin. S-functions for graphs. *Journal of Geometry*, 8(1-2):171–186, 1976.

[11] T. E. Harris. A lower bound for the critical probability in a certain percolation process. *Proc. Cambridge Philos. Soc.*, 56:13–20, 1960.

[12] H. Kesten. The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. *Comm. Math. Phys.*, 74(1):41–59, 1980.

[13] H. Kesten. Analyticity properties and power law estimates of functions in percolation theory. *J. Statist. Phys.*, 25(4):717–756, 1981.

[14] T. Kloks. *Treewidth: computations and approximations*, volume 842. Springer, 1994.

[15] T. M. Liggett, R. H. Schonmann, and A. M. Stacey. Domination by product measures. *The Annals of Probability*, 25(1):71–95, 1997.

[16] C. J. H. McDiarmid. Random channel assignment in the plane. *Random Structures Algorithms*, 22(2):187–212, 2003.

[17] C.J.H. McDiarmid and T. Müller. On the chromatic number of random geometric graphs. *Combinatorica*, 31(4):423–488, 2011.

[18] R. Meester and R. Roy. *Continuum percolation.* Number 119. Cambridge University Press, 1996.

[19] D. Mitsche and G. Perarnau. On treewidth and related parameters of random geometric graphs. Preprint. Available from `http://math.unice.fr/~dmitsche/Publications/full-treewidth.pdf`. Preliminary version in STACS 2012.

[20] T. Müller. Two-point concentration in random geometric graphs. *Combinatorica*, 28(5):529–545, 2008.

[21] T. Müller, X. Pérez-Giménez, and N. Wormald. Disjoint Hamilton cycles in the random geometric graph. *J. Graph Theory*, 68(4):299–322, 2011.

[22] M. D. Penrose. The longest edge of the random minimal spanning tree. *Ann. Appl. Probab.*, 7(2):340–361, 1997.

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

[24] M. D. Penrose. *Random Geometric Graphs.* Oxford University Press, Oxford, 2003.

[25] N. Robertson and P. Seymour. Graph minors. II. algorithmic aspects of tree-width. *Journal of algorithms*, 7(3):309–322, 1986.

Anshui Li and Tobias Müller
Mathematical Institute, Utrecht University, Utrecht, The Netherlands.
E-mails: `a.li@uu.nl, t.muller@uu.nl`.